

hp·ux/usr

Hands-On Solutions for HP-UX Users • January/February 1997



PVM—Parallel Computing on Your Desktop
SLIP on HP-UX ■ What's New with 10.x Startup and Configuration

A Publication of Interex • The International Association of Hewlett-Packard Computing Professionals

BULK RATE
US Postage
PAID
Saratoga, MS
Permit #30

GALAXY 1100

PORTABLE

HIGH-PERFORMANCE UNIX WORKSTATION



The Galaxy 1100

Complex applications demand sophisticated tools. Whether it be software development and demonstration, network administration, or other mobile HP-UX computing requirements, SAIC's Galaxy 1100 portable workstation does the job. Galaxy delivers critical information to the commercial, industrial, and military user with the performance and reliability of the HP PA-RISC processor.

Fully software compatible with the HP 9000/700 workstation series, this powerful and portable workstation offers superb graphics capability. Its 10.4-inch, full-color active matrix LCD renders images at a resolution of 1024 x 768. With minimum processing power of 79 SPECfp92 and 122 SPECint92, the Galaxy is equipped to handle even the most demanding computations.

Key Features

- 80 MHz HP 712 PA-RISC CPU
- 32 MB of RAM expandable to 128 MB with error detection/correction capability
- 2 GB hard drive
- 3.5-inch 1.44 MB floppy drive
- 10.4-inch color active matrix LCD
- 2 integral PCMCIA slots
- IEEE 802.3 Ethernet interface
- SCSI-2 interface
- POSIX-based HP-UX operating system



Science Applications International Corporation
10770 Wateridge Circle, San Diego, CA 92121
Tel 619-552-5334 or 1-800-447-4373 Fax 619-552-5253
<http://www.saic.com/it/stp/galaxy/index.html>

CIRCLE 138 ON READER SERVICE CARD

Is It New Or Is It Refurbished



**The only difference
is the money you save.**

With 17 years of knowledgeable experience and reliable service, Monterey Bay Communications is a leader in Hewlett-Packard workstation remarketing. We're professionals at providing HP 1000 and 9000 users with reliable equipment that is functionally and cosmetically identical to what HP offers — and at substantial cost savings. In addition to the 700 / 400 / 300 / 200 series, Monterey Bay Communications also offers mass storage systems,



monitors, memory and interfaces, as well as a variety of printers.

All equipment is warranted and eligible for

Hewlett-Packard maintenance. An extensive parts and spares inventory and knowledgeable staff ensure prompt service and immediate delivery.

For more detailed information or a price quotation, give MBC a call at 408/429-6144.

	NEW	MBC
Performance	✓	✓
Warranty	✓	✓
Support / Tech Expertise	✓	✓
Accessories	✓	✓
Maintenance Eligibility	✓	✓
Substantial Cost Savings		✓
Simple Order Processing		✓
Immediate Delivery		✓

CIRCLE 7 ON READER SERVICE CARD

MONTEREY BAY
C O M M U N I C A T I O N S

The HP Workstation Remarketing Specialists

Monterey Bay Communications Inc., 1010 Fair Avenue, Santa Cruz, CA 95060 Tel: 408-429-6144 Fax: 408-429-1918

Harness the Power of your Distributed Environment



PATROL

The true potential of your distributed environment is right there, waiting to be harnessed. The PATROL® Management Suite of products gives *you* the power. PATROL helps you better support your end users by ensuring their critical applications and data are available – 24 hours a day, 365 days a year.

PATROL

proactively monitors a wide range of availability and performance indicators for each managed software solution, discovering errors and correcting problems – when authorized – before end users are affected. Supporting higher availability and performance for more leading applications, databases, middleware, Internet servers, and resources than any other management product. So your end users remain productive.



PATROL management support is available today for more than 30 leading software solutions, including SAP R/3, PeopleSoft, Oracle, Sybase, Informix, TUXEDO, Lotus Notes, Internet servers, and your in-house applications. Scaling efficiently across thousands of servers running Unix, NT, OS/2 or MVS. And integrating seamlessly with your existing system and network management frameworks, including HP OpenView.

PATROL technology is being used on tens of thousands of servers around the world. PATROL is the only application and data management product shipped with every Digital AlphaServer worldwide, and PATROL will be the only management product embedded on Intel's new i960 processor. No wonder leading analysts and consultants are calling PATROL the "clear choice" for managing applications and data across distributed systems.

BMC[®]
SOFTWARE

For a copy of the most recent analyst reports and other information, CALL **800 841-2031** or 713 918-8800 or visit us on the Web at: **www.bmc.com/patrol**

Harness the Power-Today!

select countries. ® and ™ indicate USA registration or USA trademark. Oracle and Informix are registered trademarks of, respectively, Oracle Corp. and Informix Software, Inc. Other logos and product/trade names mentioned are registered trademarks

CIRCLE 23 ON READER SERVICE CARD

Contents

Features

PVM—Parallel Computing on Your Desktop

by Weicheng Jiang 30

What's New with 10.x Startup and Configuration

by John Fenwick 36

SLIP On HP-UX

by Larry Headlund 52

Departments

Product Focus

10

New Products

12

Q & A

18


HP 1000 Guru

72

CSL Perspective

76

*New Products
See page 12*



Columns

HP-UX Systems Administration

by Chris Curtin 24

HP-UX

by David L. Totsch 28

On-Site

by Judy Tarutz 60

Windows NT

by Bob Combs 64


X-Watch

by Larry Headlund 66

Internet Goodies

by Joe Berry 68

page 52



Hands-On Solutions for HP-UX Users

*Cover Story:
See page 30*



Using NFS to Connect PCs to UNIX?



We're Happy to Announce the Alternative.



Is your PC-based NFS solution providing less-than-expected performance? With TotalNET Advanced Server (TAS) software, you can share files, printers, and applications among Windows (3.x, 95, NT), NetWare, OS/2, Macintosh and UNIX computers without adding or changing anything to the client computer.

TotalNET Advanced Server delivers distinct advantages over PC-based NFS solutions:


Reduced Administration: TAS installs directly on and is administered from the server, greatly reducing time spent on LAN administration.

File Integrity: TAS offers precise file and record locking for all network operating systems it supports.

Protocol Independence: TAS includes NetBEUI, NetBIOS, TCP/IP, IPX/SPX, and AppleTalk protocols for compatibility with most native client computers.

No Additional User Training: TAS emulates all popular PC LANs, enabling users to connect to the server and share resources in each client's own environment.

Test drive the TotalNET answer to high-performance networking. Call or visit our website today for a free network evaluation.



SYNTAX

2 0 6 8 3 8 2 6 2 6
<http://www.syntax.com>

© 1996 Syntax, Inc. All Rights Reserved.
All trademarks of companies whose products are referenced are hereby recognized.

CIRCLE 163 ON READER SERVICE CARD

<http://www.interex.org/>

INTEREX EXECUTIVE DIRECTOR
Charles A. Piercey

PUBLICATIONS MANAGER
Connie Wright

EDITORIAL
MANAGING EDITOR
Michael Ehrhardt

COPY EDITOR
Richard Kranz

NEW PRODUCTS EDITOR
Michelle Pollace

PROOFREADER
Jean Nattkemper

ADVERTISING & MARKETING

ADVERTISING SALES MANAGER
Brian Hallin

ADVERTISING REPRESENTATIVE
Kathie Schwartz

ADVERTISING COORDINATOR
Phil Nguyen

DESIGN AND PRODUCTION

SENIOR GRAPHICS DESIGNER
Molly McGinnity

DESKTOP PUBLISHING SPECIALIST
June Ramirez

GRAPHICS PRODUCTION SPECIALIST
Gale Patterson

hp-ux/usr is published bimonthly by Interex, the International Association of Hewlett-Packard Computing Professionals. Second-class application pending at Sunnyvale, California 94086 and additional offices. The editorial and business offices are located at 1192 Borregas Ave., Sunnyvale, California 94089, USA, 408.747.0227, Fax 408.747.0947. Address membership questions and change of address to Membership Services. Address all questions concerning circulation/distribution to the Distribution Manager.

Remittances should be sent to Interex,
File No. 61054, P. O. Box 60000,
San Francisco, California 94160, USA.

Address all editorial correspondence to Michael Ehrhardt,
Editor, *hp-ux/usr* Magazine, c/o Interex,
P.O. Box 3439, Sunnyvale, California 94088-3439, USA.

Subscription to *hp-ux/usr* is \$49.50 (6 issues) per year in the U.S., add \$25.00 for Canada and Mexico, add \$50.00 for all other countries. Member Services (Associate, Contributing or Online Service Package membership levels) include a subscription to *hp-ux/usr* at \$49.50. For other Member Services refer to membership form.

Statements of fact and opinion are the responsibility of the authors alone and do not imply an opinion on the part of the Interex Board or Magazine. Entire contents copyright © 1997 by Interex. All rights reserved.

POSTMASTER: Send address changes to:
Interex, P.O. Box 3439,
Sunnyvale, California 94088-3439 USA.
Attention: Member Services.

TRADEMARKS: UNIX, X/Open Co. Ltd.,
HP-UX, Hewlett-Packard; X Window System,
X Consortium, Inc.

How to Contact Interex...

EDITORIAL SUBMISSIONS

hp-ux/usr encourages readers to contribute their opinions, tips, and solutions. When sending letters for publication or to request author contribution guidelines, please address them to *hp-ux/usr* editor Michael Ehrhardt.

Postal Address:
Interex
P.O. Box 3439
Sunnyvale, CA 94088-3439

Office Address:
1192 Borregas Avenue
Sunnyvale, CA 94089

Because of the difference in zip codes between our office address and P.O. Box, please be sure to address all regular mail to the P.O. Box. Any express service packages should be delivered to the Borregas Avenue address. Thank you for your attention to this small but significant detail.

TELEPHONE:

The Interex switchboard is open 8:00 a.m.–5:00 p.m., Pacific Time. Call 800.468.3739 (U.S. and Canada) or 408.747.0227. After 5:30 p.m. our voicemail system will record your call.

FACSIMILE:

Call 408.747.0947

INTERNET:

To send e-mail to Interex, use the following address format:

<IDname>@interex.org

The following IDs are currently active on Interex's HP 9000:

Address	Department
ehrhardt	<i>hp-ux/usr</i> Letters to the Editor, Q&A, and requests for author guidelines
webmaster	Internet support
csllhpux	Contributed Software Library
membership	Membership/subscription inquiries and services
pubs	Circulation and advertising inquiries
conference	Conference questions and arrangements

An example would be `pubs@interex.org`
Anything before the @ sign is case insensitive.

COMPUERVE:

Interex can be contacted via the CompuServe ↔ Internet gateway. To send CompuServe mail, use the following format:

>INTERNET: `pubs@interex.org`

You can address your mail to specific departments using the ID's listed above.

Interex maintains a CompuServe account that is collected daily. Please address all messages to ID no. 76376, 1222.



Shared Knowledge.
Shared Power.

As a not-for-profit association of HP computing professionals, Interex is dedicated to meeting the information, education, and advocacy needs of its members worldwide.

In today's world of rapidly changing technology, Interex puts hands-on solutions to hardware, software, and operating system problems at your fingertips. Because members actively contribute—exchanging ideas and sharing solutions—Interex is a vital link in the transfer of HP expertise.

Operating independently from Hewlett-Packard, Interex has more than 20 years of serving HP computing professionals. Through its publications, conferences, and volunteer committee structures, Interex has the qualifications to represent you, a valuable member of the HP user community.

Interex® is a trademark registered in the U.S. Patent and trademark office.

your most valuable asset

and how to cover it.

Corporations protect their people with comprehensive health insurance.

They protect their buildings and capital equipment with catastrophic insurance. But what about the asset that many professionals feel is their most valuable? Data.

We are Concorde Technologies. And our network backup solutions are really data insurance. Case in point? Our 2/10 DLT® Tape Systems. Perfect for midrange applications, these high performance, 200-gigabyte libraries house two DLT 4000 tape drives and up to 10 DLT tapes. Their dual access technology optimizes system speed and delivers "lights out" backup, even if a drive goes down. An intuitive front panel with a diagnostic display gives you the status of all important library operations, instantly. And because you can connect up to nine 2/10's together where all drives can access all tapes, Concorde Technologies gives you a solution that meets your needs today—and tomorrow.

Bottom line, we give you an edge that goes beyond any specific boxes we sell. That edge is knowledge.



We know heterogeneous network backup as no other HP Channel Partner does. We know the right process, products and service. We know how to analyze your requirements and design the optimal solutions. And we know how to install, train and support the system.



In short, when it comes to your company's data, your asset is covered with Concorde Technologies.

Solutions beyond the box.

800-359-0282 • (619) 536-5500 • Fax (619) 566-4396

email: sales@concordetech.com

URL: <http://www.concordetech.com>





Editor's Notes

Regular readers of *hp-ux/usr* will have noticed Bob Combs's Windows NT column. We added the NT column to the magazine with the September/October issue, in response to the ever-growing presence of NT workstations and servers within HP-UX systems. In coming issues we will have features on HP-UX and Windows NT integration and interoperability.

Although the primary focus of *hp-ux/usr* remains the HP-UX user, both technical and commercial, we cannot ignore the trend toward mixed shops that have inter-operating computer systems from a variety of different vendors. The magazine's mission is to be a solutions provider and, as such, it must embrace the widening range of technologies HP computing professionals have at their disposal.

This issue's cover story, by Weicheng Jiang, is an introduction to PVM—Parallel Virtual Machine, software that enables you to write portable parallel programs that can be executed across multiple nodes on a network. If you're crunching the numbers fast and furiously, here's a method you will find useful. High-speed network technologies such as Fibre Channel make it possible for you to build your own parallel computers, harnessing the combined power of a loose group of workstations.

Larry Headlund, our X-Watch columnist, provides the lowdown on SLIP with HP-UX. This poor relation of networking, implemented with serial cabling and modems, is a choice only when you have nothing better, but you may find yourself connecting to a machine via SLIP. Since SLIP isn't the method of choice if a faster connection is available, the initial installation of HP-UX doesn't set everything up for you to use SLIP. Headlund's article describes the procedures for setting up SLIP on an HP-UX system.

John Fenwick played to packed houses at the most recent InterWorks and HP World conferences. And what was the play about? Startup and configuration in HP-UX 10.x, of course. With HP-UX 10.x comes not only a "normal" UNIX file system layout, but a new way of controlling the subsystems. 10.x startup and configuration resembles designs used by other major UNIX vendors. Gone are the large */etc/*rc** files: these have been split up into functional units. System administrators and developers now control system startup and shutdown by modifying configuration files. Execution scripts have been separated from the configuration information. Fenwick's article is a thorough introduction to HP-UX 10.x startup and configuration—system administrators and developers on the upgrade path to 10.x will want to keep it handy for future reference.

M. M. Ehrhardt

Michael Ehrhardt
Managing Editor



LOOKING FOR A FASTER WAY TO LOAD YOUR DATABASE?

SyncSort UNIX can help you complete database loads, reorgs and reports in as little as *half the time*.

SyncSort combines high-speed sorting, versatile data manipulation features and the ability to handle a variety of data and file types. The result is a powerful, flexible tool for breaking database bottlenecks. To order SyncSort or for a free copy of our booklet, "Sorting and Relational Database Performance," please call or fax.

syncsort

Tel (201) 930-8200 dept. 17HUS
Fax (201) 930-8285 dept. 17HUS



Product Focus

Altia Design/ProtoPlay

For most companies, establishing a presence on the Web is a lot like devising an elaborate Yellow Pages ad. Prospective clients can run a search by concept or company name to arrive at that company's home page, which typically contains information in a text-and-graphic format akin to that which is represented on a printed page.

To increase the usefulness of the Web as a marketing and collaboration tool, Altia, Inc. has devised ProtoPlay, a plug-in module to the Altia Design 2.0 prototype designer. Using a PC-compatible computer and standard modem speeds, Internet users can access an electronic prototype and "play" with products—such as cellular phones—as if using the real thing.

Michael Juran, president of Altia, Inc., said ProtoPlay allows users to convert prototypes developed with Altia Design into "some useful and valuable content for the Web"—content that is more sophisticated than the text and graphics of printed material and that exceeds the sophistication of multimedia Web presentations. ProtoPlay, Juran explained, exploits capabilities truly "unique to computers and the Internet"—interactivity.


While interactivity is implemented on the Web in varying degrees, Juran observed that "interactive" describes anything from fill-in forms to virtual reality, which today "is more entertainment than useful value," especially when implemented over the Internet with bandwidth restrictions.

Juran believes the timing of Altia's technology and that of the Web have converged; extending Altia Design prototypes to the Web is a natural evolution, because the application was designed to generate prototypes small enough to fit on a floppy disk, be sent through e-mail, and be loaded into laptops remotely. Using ProtoPlay, even nonprogrammers can place small prototype files on their Web sites. The resulting prototypes have all the functionality of actual products, Juran said, and they can be deployed "on the Internet as it is today"; such prototypes can be loaded into a

PC over a 14.4 modem with ease. (Minimum system requirements for Design/ProtoPlay are a 33-MHz 486 PC with 8 MB RAM. UNIX systems run the applications easily, so no minimum requirement is specified.)

Juran believes Web-based prototypes exploit the relatively unexplored capability of the Web to create markets and interest for products through interaction. His vision will be put to reality soon enough, as "very shortly we will see some companies that have ProtoPlay on the Internet and that are using it as a sales tool." While not able to name the companies at the time of this writing, Juran did reveal that a few client companies might have their

An electronic prototype of a cellular phone, which allows users to play with the features on the Web



sales tools up by the end of 1996 and that Altia would divulge more information at that time.

Today many companies, including HP, use Altia Design/ProtoPlay over intranets "as a way to do collaborative design." Used in this capacity, ProtoPlay enables designers to draw upon another one of the Web's inherent strengths: version control. Because the prototype is placed in one location on the Web site, it requires collaborators to access the same prototype, eliminating the need to track different prototype versions.

Just as the generated prototypes are designed to be easy to implement, Altia Design and ProtoPlay are designed to be easy to use. "The gap we try to bridge is [that between] programmers and less technical people, such as those in marketing or human factors," Juran said.

Altia Design implements a design method Juran calls "direct manipulation," by which the designer can use a mouse to move a design object to another location on the screen and instruct the program to store that movement as one "state" in a sequence; for example, a movement identified as "state 1" would be the first action to occur for a given sequence triggered by a user's interaction with the prototype. Designers can specify any number of states for a given interactive stimulus from the user, and the program will interpolate, or "fill in," actions for the prototype to take wherever gaps exist between programmed states.

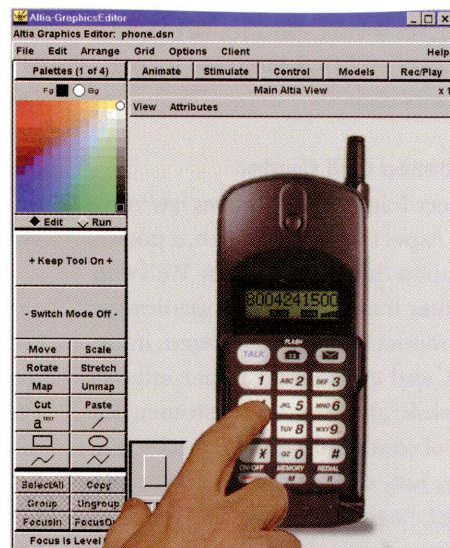
Designers can also specify "stimuli" for the prototype. For example, to program the stimulus for a knob, the designer would draw a circle (or knob), rotate the knob through direct manipulation using the mouse, then specify the states a left-mouse click represents, and so on. Stimuli specification is achieved through series of dialogue boxes, and designers can also program prototypes to accommodate touch screen as well as mouse input.

Designers can also add logic to prototypes, but as with programming movement and stimuli, no coding is required. Altia Design can also connect to C, C++, and Visual Basic code as well as to other modeling and simulation tools.

Because the product is meant to be used by nontechnical and technical product designers, support of both PC and UNIX platforms is essential. Juran explained, "In engineering the primary design platform is still UNIX, and in marketing and human factors disciplines, people use PCs." Altia Design and ProtoPlay run under Windows 3.1, NT, or 95 and on UNIX workstations by HP and others.


A complete Altia Design system costs \$5,900 for the PC and \$9,900 for UNIX workstations. ProtoPlay is a free plug-in to Altia Design 2.0 and is available on Altia's Web site. To enable users to widely distribute prototypes, Altia assesses no runtime royalty fees.

Contact Altia Design, phone: (719) 598-4299, fax: (719) 598-4392, <http://www.altia.com>. ■



**Altia Design editor
used to build
prototypes. It
requires no
programming.**

Michelle Pollace, the New Products editor for hp-ux/usr magazine, writes Product Focus.



New Products

ERP Manufacturing Solution

Symix has announced Symix SyteLine, an advanced client-server solution for mid-sized manufacturers of discrete, configurable products. Symix SyteLine Version 2 builds upon The SYMIX Solution, the company's flagship Enterprise Resource Planning (ERP) product. The newly designed customer service component takes advantage of a new event-driven GUI to allow faster, more flexible navigation through the system. Enhanced multisite capabilities support both centralized and decentralized order entry structures, providing full visibility into inventory at all sites at all times. Companies that require sophisticated financials can better report or consolidate information on current or cumulative fiscal status.

Designed to operate in a client-server environment, SyteLine is written in PROGRESS and supports multiple platforms, including Windows NT and UNIX. New licenses range from \$100,000 to over \$1 million, depending on consulting and services offerings and the number of concurrent sessions.

Contact Symix, phone: (614) 523-7000, fax: (614) 895-2504, <http://www.symix.com>.

Web-based Call Center

Aspect Telecommunications has introduced Aspect Interactive Web, a powerful solution that integrates the Web with call center transactions. Companies can now conduct more personalized, informative, and efficient customer interactions while allowing the customer full choice of communication method—telephone, fax, electronic mail over the Internet, pager, and now, the Web.

Users can create Web pages tailored to individual customers, call back Web cus-

tomers who request live interaction, automatically notify customers of important information (such as a change in stock price), and produce comprehensive reports covering all customer interactions, regardless of access method. Aspect Interactive Web is featured in a new release of Agility, Aspect's interactive response system, which delivers automated customer service solutions in a call center environment. Agility 2.0 also offers speech recognition and text-to-speech capabilities.

Pricing for Agility 2.0 begins at \$75,000. The optional Aspect Interactive Web license starts at \$36,000. The speech recognition and text-to-speech options start at \$7,500 and \$600, respectively. Agility 2.0 can handle from 2,000 to 12,000 transactions per hour.

Contact Aspect Telecommunications, phone: (408) 325-2483, <http://www.aspect.com>.

Free Computing Dictionary

O'Reilly & Associates' new *Dictionary of PC Hardware and Data Communications Terms* is now freely available on the Web at <http://www.ora.com/reference/dictionary/>. O'Reilly created this online version of the *Dictionary* to give Web users and content providers a convenient way to look up definitions "on the fly." Each of the more than 900 terms has its own URL, and the Web site includes a search engine.

Author Mitch Shnier is keeping the *Dictionary* current on the Web so that Web content developers can link to the terms from their documents. The *Dictionary of PC Hardware and Data Communications Terms* offers more information than a typical dictionary. For example, the description of PGP (Pretty Good Privacy), a popular data encryption protocol, covers three and one-half pages. Terms are care-

InterWorks '97

Conference & Exhibition

- System Administration Conference
- Software Development Conference
- Internet / Intranet Training
- Windows NT Highlights

APRIL 12 – 17, 1997
PHILADELPHIA, PA

<http://www.InterWorks.org>

Software Tracking

Wyatt River Software (formerly Central Design Systems, Inc.) has announced LicenseTrack 4.0, a powerful, scalable tool enabling software license management across LAN, WAN, and enterprise networks. LicenseTrack 4.0 gives IS managers and system administrators the tools to quickly evaluate software usage trends and status. It runs on DOS, Windows (3.x, 95, NT), and all major UNIX operating systems.

LicenseTrack 4.0 now features Agents for Windows 95 and Windows NT, allowing the tracking of 16- and 32-bit applications and Java applets; a GUI-based console for managing FLEXlm daemons; integration with systems management tools; significantly enhanced reporting capabilities; improved utilities; and easier installation. LicenseTrack 4.0 pricing starts at \$995 for PC servers and \$4,995 for UNIX systems.

Contact Wyatt River Software, phone: (408) 327-8840, fax: (408) 327-8844, <http://www.wyattriver.com>.



Wyatt River Software, LicenseTrack 4.0

management of mission-critical Windows NT servers and workstations.

The remote power manager enables administrators to initiate an orderly Windows NT shutdown for a remote unattended server by using the Windows NT UPS Service low-power monitoring function. After Windows NT is safely shut down, the product's multilevel password interface allows the user to turn on or off the remote system securely over a communications line.

Sentry ShutDown interacts with any standard ANSI terminal emulation program. It supports out-of-band modem or RS-232 communications; an optional model supports an in-band telnet TCP/IP session.

Sentry ShutDown Remote Power Manager is priced at \$669.95 for 110-VAC power modules and shutdown signal cable, or \$719.95 for 230-VAC power modules and shutdown signal cable.

Contact Server Technology, phone: (800) 835-1515 or (408) 745-0300, <http://www.servertech.com>.

Factory-Wide Systems GUI

FASTech Integration, Inc. has announced the Intelligent Windows NT Client, which is designed to provide integrated equipment control, statistical analysis, production reporting, and desktop integration with a common GUI to factory-wide systems. FACTORYworks 2.0 includes the Intelligent Windows NT Client, which is fully ODBC-compliant and based on industry standards. FASTech's NT Client's ActiveX controls, combined with the Visual Basic environment, creates a rapid business rule development engine, with which the

technique intended to provide a more accurate picture of the costs relating to the manufacture and sale of goods. The Quality Management module was added to help companies improve manufacturing efficiency, comply with customer quality standards and regulatory requirements, and minimize rework and rejects. The Change Management module links to other applications within System 21 to ensure that changes to the manufacturing process have been properly authorized and that they adhere to regulatory or customer standards.

Contact JBA International, phone: (800) JBA-INTL, ext. 3033, fax: (609) 439-9652.

Remote Power Control

Server Technology, Inc. has introduced Sentry ShutDown Remote Power Manager, which provides a secure multilevel password interface for power

fully cross-referenced.

The printed book (*Dictionary of PC Hardware and Data Communications Terms*, 532 pages, ISBN: 1-56592-158-5, \$19.95) is available from O'Reilly and in bookstores worldwide.

Contact O'Reilly & Associates, phone: (800) 998-9938 or (707) 829-0515 fax: (707) 829-0104.

Manufacturing Solution

JBA International has unveiled three new modules for its client-server-based System 21 Manufacturing software: Activity Based Costing, Change Management, and Quality Management. System 21 Manufacturing is part of JBA's System 21 software suite, which also includes packages for managing financials, customer service, logistics, and service management. All operate on UNIX and AS/400 platforms.

Activity Based Costing is an accounting

business rule writer can integrate highly complex manufacturing scenarios.

Users can easily navigate through the factory site model, zooming in to access specific lot, equipment, or attribute data, or zooming out to view all the lots for a particular bay.

With a few lines of Visual Basic code, users can execute site-specific actions such as download an equipment program, collect equipment data, activate a user input window, or initiate a transaction to an external system.

Contact FASTech Integration, phone: (617) 259-3131, fax: (617) 259-3188.

Web Agent

Vital Inc. has announced Trudger, which makes Web surfing possible without a live connection. A preview version can be downloaded from Vital's Web site for multiple UNIX and Windows platforms.

Trudger allows users to download Web pages, graphics, text, and HTML links of interest. Trudger brings back all updates automatically. Users can configure Trudger to control the type of content/files to download, number of links to traverse, Web sites to avoid, time to spend downloading, amount of disk space to use, etc.

Trudger can stop an in-progress download and resume it later. Users can specify how many links to download in parallel and adjust the transfer rate while it is in operation. Users can also pause an in-progress download and then delete the links/pages they do not want to download.

Trudger does not need a Web browser to perform its function, but it can communicate with one to display the downloaded pages.

Trudger is currently available on HP-UX and other platforms and is priced

at \$59.99.

Contact Vital Inc., phone: (214) 612-2684, fax: (214) 612-3326, e-mail: info@vital.com, <http://www.vital.com>.

COBOL-to-ODBC Interface

Acucobol, Inc. has announced Acu4GL for ODBC, said to be the first product to interface seamlessly from COBOL to ODBC-compliant data sources. Acucobol COBOL applications can now access popular databases such as Oracle, Access, and Informix. Because the Acu4GL interface takes advantage of standard ODBC drivers, developers can communicate with many different database formats without having in-depth knowledge about them, and with-

out writing customer database queries.

Applications need to be compiled only once to instantly plug into any ODBC-compliant data source. Built into Acu4GL's sophisticated interface is a complete understanding of the language and standards of both COBOL and ODBC. Acu4GL works behind the scenes to satisfy each request of the executing COBOL application by generating the appropriate ODBC calls.


Contact Acucobol, phone: (800) COBOL-85 or (619) 689-7220, fax: (619) 566-3071, e-mail: info@acucobol.com.

Web Load Balancing

HydraWEB Technologies has announced HydraWEB, a new patent-

Advanced Geometric Functions

Building Block Software has announced the CCDK/Advanced Geometry developer's kit, with a set of over 150 advanced geometry C functions. It supports operations with areas, regions, and polylines. This type of graphics processing is becoming increasingly important to Internet Web page authoring tools that define arbitrarily-shaped regions for clicking and painting.



CCDK/Advanced Geometry

The CCDK/Advanced Geometry developer's kit is compatible with DOS, Windows (NT/95/3.1), and UNIX operating systems. It requires one of the following compilers: Microsoft Visual C/C++ 1.x/2.x/4.x, Borland C/C++ 3.x/4.x/5.x, WATCOM 9.x/10/x, or UNIX C/C++ compiler (HP and others).

Contact Building Block Software, phone: (617) 860-9091, fax: (617) 860-9066, e-mail: info@buildingblock.com.

Prolifics, next-generation application development environment

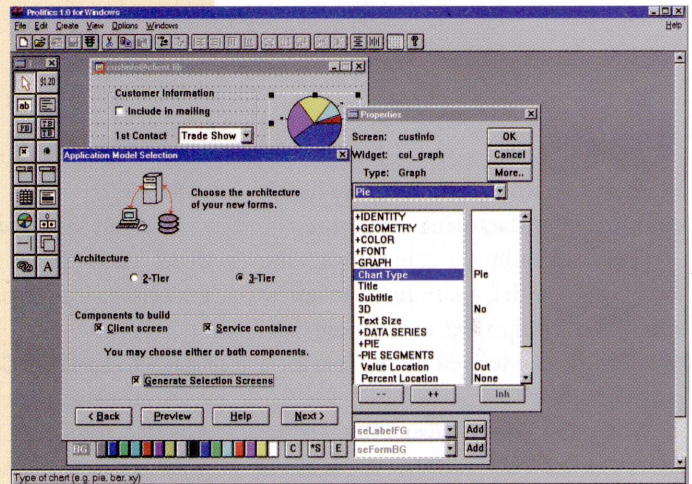
Enterprise Transactional Applications

Prolifics has announced a next-generation application development environment that lets users build, deploy, and maintain enterprise-class database transactional applications both on the Web and in 3-tier partitioned client-server environments. Prolifics combines the industry-standard, high-performance TUXEDO transaction processing middleware with Visual Server Development.

Prolifics' embedded OLTP middleware, JetNet, is based on BEA TUXEDO, a proven, open middleware solution. JetNet provides distributed application processing across a back-end network of servers, enabling superior application performance and scalable growth.

The Prolifics development environment is priced starting at \$35,000 for a five-user development system.

Contact Prolifics, phone: (212) 267-7722, fax: (212) 608-6753, <http://www.prolifics.com>.



pending load-balancing system that multiplexes Web servers to ensure 24-by-7 Internet/intranet operations. HydraWEB provides a fault-tolerant Web system that allows system administrators to load-balance HTTP requests across multiple servers or clusters. Serving as the link between Web servers and the Internet or intranet, HydraWEB manages all HTTP requests and intelligently forwards them to the best available server.

If a Web server goes offline or experiences hardware or software failure, HydraWEB automatically rebalances the load among the remaining servers, in real time. Fully supporting unlimited Web servers, HydraWEB ensures 24-by-7 operations in addition to offering a high-performance server cluster to accommodate large numbers of "hits." It optimally uses heterogeneous machines of different sizes, capacities, and hardware and software configurations. HydraWEB supports over 1 million hits per hour.

Each server managed by HydraWEB can be taken offline without shutting down or disrupting the entire Web site.

HydraWEB is priced per server and ranges from \$4,500 for a two-server license to \$30,000 for an unlimited server license.

Contact HydraWEB Technologies, phone: (908) 972-5252, e-mail: pr@hydraweb.com.

Order Management

Industri-Matematik International Corporation (IMI) has announced a new version of System ESS software, which provides manufacturers and distributors with new capabilities to meet the diverse order handling requirements of customers and external partners. The software also provides an Internet component for accessing customer order information using the Web.

Included among the new System ESS features are expanded order handling capabilities to help manufacturers better support the popular industry initiative, Continuous Product Replenishment; enhanced pricing and promotions management for consumer packaged goods and industrial products companies; real-time monitoring and alerts of logistics activities; and enhanced connectivity to legacy systems.

Designed to work with standard Web browsers, the new System ESS Order Tracker enables companies to extend internal applications by offering customers and field sales organizations

access to shipping and order information through the Internet and intranets.

System ESS 4.3 is priced on user configuration and typically starts at \$600,000. The solution runs on HP-UX and other open systems platforms.

Contact IMI, phone: (914) 631-2700, fax: (914) 631-5111, <http://www.imi.se>.

Supply Chain Optimization

Red Pepper Software has plans to deliver ActiveAgent object messaging integration between Red Pepper's ResponseAgent supply chain optimization solutions and SAP R/3. The end solution will serve to integrate all of Red Pepper's ResponseAgent products with logistics modules in SAP R/3 Version 3.0.

Red Pepper's ActiveAgent real-time integration leverages SAP's Workflow and ALE technologies. For example, as new customer orders are processed in R/3, an R/3 Intermediate Document is generated and published as a workflow event. A corresponding ActiveAgent subscribes to this message and provides it to the appropriate ResponseAgents, directly updating their memory supply chain models with the latest customer orders. If the ActiveAgent is configured to monitor all sales order, this integration is truly real-time.

Contact Red Pepper Software, phone: (415) 525-3300, fax: (415) 341-8064, <http://www.pepper.com>.

Continued on Page 78

**Quick: Name the RAID system picked #1
by PC Magazine, BYTE and LAN Magazine.**

StorageWorks™

When industry's toughest independent critics put RAID subsystems to the test, they came to a unanimous conclusion. StorageWorks™ RAID subsystems from Digital™ delivered outstanding availability, reliability, and economy—plus industry-leading performance. All designed as scalable solutions for today's multivendor PC and UNIX networks, to protect your investment from desktop to datacenter.

To find out more about Digital StorageWorks RAID solutions, call Workstation Technologies at (800) 663-1966.



Workstation
Technologies, Inc.

"Solutions for Your Open Systems Computing Environment"

1-800-663-1966

Outside the U.S.: 603-890-6700

Digital, the Digital logo, StorageWorks and the StorageWorks logo are trademarks of Digital Equipment Corporation. Other trademarks are property of their respective holders.

CIRCLE 20 ON READER SERVICE CARD



Question & Answer

Q: When I click on the terminal on the front panel of my HP Vue screen, the hpterm window that is started does not read the login shell scripts such as *.profile*. How can I change this?

A: This is the default behavior for HP Vue and CDE. To change this behavior you must first decide if you wish to make a global change for all users on this system, or a local change just for your user. The steps to do both are defined below.

Remember that HP Vue is highly customizable. There are many different ways to achieve the same result. You may choose to use a procedure different from the ones described below and still get the desired results.

To make this change only for your own user follow these steps:

Method 1 (for VUE, one user at a time)

1. First check to see if your *\$HOME/.vue* directory contains a *vuewmrc* file. If it does, go to step 2. If it does not, you must copy the file as follows. Be sure to change the copy's properties to grant you write permission:

```
cp /usr/lib/X11/vue/Vuewm/sys.vuewmrc $HOME/.vue/vuewmrc
```

2. The action that gets executed when you click on the terminal control of your HP Vue front panel is defined in the *vuewmrc* file as *TERMINAL*. *TERMINAL* will not execute your login scripts but *HPTERM_LOCAL* will. Change *all* occurrences of *TERMINAL* to *HPTERM_LOCAL* in this file:

```
Terminal [P] @term.m.bm f.action TERMINAL
```

After the change:

```
Terminal [P] @term.m.bm f.action HPTERM_LOCAL
```

3. The only remaining step is to set the resource that will cause *HPTERM_LOCAL* to execute the login shell scripts.

The resource is the following:

```
hpterm*loginShell: True
xterm*loginShell: True
```

The simplest way is to type these lines exactly as they appear above into a file and then merge that file into your resource database using the following command:

```
/usr/bin/X11/xrdb -m yourfilename
```


Using the RES_EDIT action defined in your file manager is another way to set this resource.

4. Finally, close all hpterm windows except the console and then log out of HP Vue and log back in. This is necessary to force HP Vue to reread the *vuewmrc* file and to kill all existing hpterm windows and cached hpterm windows.

If you now click on the terminal on the front panel to start an hpterm window, your login shell scripts will be read. It is very important to note, however, that while the changes to your *vuewmrc* file are permanent, the setting of the *hpterm*loginShell* resource will survive only if you return to your current session or if you save your current session as your home session and return to your home session.

Method 2

Make this a global change for all users on this system:

1. If the users have their own *\$HOME/.vue/vuewmrc* files, they must be changed the same way as in step 1 of the first procedure described above. If a user does not have this file, the same changes must be made to the */usr/lib/X11/vue/Vuewm/sys.vuewmrc* file. This will then affect all users without their own copy of this file.
2. Again the *hpterm*loginShell* resource must be set to true. To make this global for all users, one option is to create a file called HPterm in the directory */usr/lib/X11/app-defaults* directory. This file would then contain the resource name as follows:

```
hpterm*loginShell: True
```

For the file *usr/lib/X11/app-defaults/Xterm*:

```
xterm*loginShell: True
```

For the file *usr/lib/X11/app-defaults/Dterm*:

```
dterm*loginShell: True
```

This file will be read each time a terminal emulator is started and thus it will read the resource and in turn your login shell script.

3. Close all hpterm windows except the console and log out of HP Vue and back in. This will force HP Vue to reread

your *vuewmrc* file and restart term windows and cached term windows.

These changes will force all term windows, whether started from the front panel, the file manager, or the command line, to read the login shell scripts by default for every user.

There is a downside for */usr/lib/X11/app-defaults*: Any patches or replacement resource files that come with an upgrade will overlay these resource files. So evaluate this method as opposed to some of the other methods.

Method 3

Use the environment variable *XENVIRONMENT* and set it equal to whatever file you would like to use. For a global setting, you could put the file in */etc*, as in */etc/Xstuff* and set *XENVIRONMENT=/etc/Xstuff*. Local users can set this value in *.profile* to match their own preferences.

Method 4

For CDE, the best method is to edit */etc/dt/config/C/sys.resources*. If the file does not exist, copy the template from */usr/dt/config/C/sys.resources* into */etc/dt/config/C/sys.resources* and add the individual resources there.

Method 5

You can always use *xrdb* to merge new resources into your environment. *xrdb -m <some-file>* will add whatever additional resources you need. Methods 3 and 5 are X-based so the choice of window manager (mwm, VUE, CDE, etc.) is not important.


Q: What changes can I make to improve fbackup performance with a DDS drive?

A: The fbackup man page lists the parameters that can be set in the configuration file and their default values. However, some of the defaults are out of date now that the DDS drives are replacing half-inch magtapes and cartridge tape drives. The parameters were not designed for the second generation of DDS drives, such as the HP 2000S (non-compression), HP 2000DC, and HP 4000DC.

Before looking at the configurable parameters of fbackup, it will be useful to check some of the basic concepts relating to DDS tapes. Data is mapped onto DDS tape in groups. Each group is 128 KB in length and at the end of the group there

is a group information table (GIT) and a block access table (BAT), which is a type of index.

DDS tape group



The GIT is a fixed length of 32 bytes. The BAT is variable in length and information about its length is stored in the GIT. The data area will vary in size, becoming smaller as the BAT gets larger.

Every entry in the BAT is 4 bytes in length and there is an entry in the BAT for each file marker, fast search marker, and block in the group.

The overhead will mount quickly when the backup consists of many small files. For example, if all the files in a backup were only 10 bytes, the overhead of a 4-byte entry in the BAT for every file would amount to 40 percent of the available space for each group. By comparison, if the backup consists of 10-KB files, the overhead of space in each group drops to 0.07 percent.

It should also be noted that the DDS drive has a five-second window in which data must be flowing. If the flow of data ceases for five seconds, the DDS group (128K bytes) will be written, even if it contains only a few bytes of data. This is why data logging (data is written once every few minutes) causes very low tape capacity (as low as 100 megs!). If the system is very busy and fbackup has a poor priority, similar loss in capacity might be seen. Moral: Perform backups on quiescent or single user systems.

Here is a brief summary of the overhead incurred by fbackup:

The start of each tape has an ANSI Standard Label area of one KB and a volume Header of two KB. Following that is the file index, which varies in length according to the number of files backed up and the complexity of the directory structures.

Each data file has an overhead of a 1-KB header and a 1-KB trailer. In addition, the data is rounded up to the next 1-KB boundary, which may waste a small amount of space.

The fbackup configuration file parameters and defaults are:

blockperrecord	16
records	32
checkpointfreq	32
readerprocesses	2 (maximum of 6)
maxretries	5
retrylimit	5000000
maxvoluses	100
chgvol	/usr/adm/fbackupfiles/chgvol
error	/usr/adm/fbackupfiles/error
filesperfsm	200

The fbackup configuration file option is called with `-c <filename>`. For example:

```
/etc/fbackup -f /dev/rmt/0hc -v -i /users -c /fb_config
```

If you create a configuration file, you can set the parameters for fbackup to non-default values. You need only list the parameters that you want to change and the others will retain the defaults.

SAM uses fbackup and has its own fbackup configuration file: `/usr/sam/config/br/fbackup_config`. The parameters in this file can be adjusted too. Do not delete any that are listed, just add parameters that are not listed or alter the values of the ones that are present.

The parameter that has been shown to have the biggest impact on the performance of the compression drives is *blockperrecord*.

For most systems, adjusting this parameter to 128 will cause a noticeable improvement in both tape capacity and the time it takes fbackup to finish on a DDS drive. Some users have tested fbackup with higher values such as 256, 512, and 1024. These are all legal values for the newer drives at HP-UX 9.0x and 10.xx, but usually 128 will result in improved performance.

Some users have also increased the value of *checkpointfreq* to a value of 256. Used in conjunction with the higher values for *blockperrecord*, this parameter improves tape capacity but, if used alone, its effect will usually be marginal. The overhead of each checkpoint record is 62 bytes, preceded by an EOF marker. This overhead is within the data area of the group.

The *filesperfsm* parameter controls the number of files between each fast search mark and has an impact on the time to recover files from a backup tape. During a recovery, the tape streams rapidly to the nearest fast search mark, then slows down to check each file, up to the required one.

Setting *filesperfsm* to a low value means that there will be fewer files between each fast search mark. With older DDS drives this could result in much faster recovery of individual files. With the newer DDS drives, such as the compression drives, recovery operations are much faster than before and reducing the value of this parameter is usually not necessary.

The documented minimum value for *filesperfsm* is 2 and overhead is 4 bytes per marker. If you back up lots of very small files, and also set the *filesperfsm* value very low, enough overhead will accumulate to be noticeable. However, the difference in tape capacity, for a typical full backup with thousands of files, will amount to a few megabytes.

viNOT



THE UNIX TEXT EDITOR FOR NORMAL PEOPLE

- **Easy to use**
- **Easy to customize**
- **Easy on your budget**

For an evaluation copy or additional information, contact:



Computer Solutions, Inc.
of Orlando

120 E. Marks St. • Suite 225 • Orlando, FL 32803
407-649-0123 or 512-343-6634
FAX 407-649-1407 email - info@csisoft.com

CIRCLE 153 ON READER SERVICE CARD

Fast search markers can affect backup times. The SCSI driver interrupts the normal writing flow to create fast search markers and this can take up to four seconds for each one. It's probably best not to alter the default for *filesperfs* unless your recovery operations on individual files are really taking a long time. On average, it should take between one and three or four minutes to recover a single file when *filesperfs* is optimal.

The actual performance values achieved with the fback-up configuration file will vary in different circumstances. It is worthwhile for system administrators with a DDS drive to test several variations of these parameters with their own backups. The *blocksperrecord* parameter is the best one to experiment with first. Remember to adjust values one parameter at a time and keep records of the actual time. (Hint: see the *time(1)* command.)

Q: I have a 735 workstation and an 816 business server. Both are running 10.01. Is it possible for me to compile my software on one system and run it on the other?

A: You should be able to execute on 700 systems at 10.01 code that was compiled on 800 systems at 10.01. But code that

was generated on 700 systems would need to be compiled with

+DA1.0 and *+DS1.0*

Q: I have two Series 700 workstations. One is at 9.05 and the other is running 10.01. I am transferring files from one system to the other via ftp(1). In the past, when transferring between 9.x systems, I've used *cksum(1)* to verify the transfer. This doesn't seem to work with the different levels of the operating system. The checksum values don't match.

A: The *cksum(1)* command at 10.01 is using a different algorithm to conform to the POSIX.2 standard. You may use the command

```
/usr/old/usr/bin/cksum
```

on 10.x systems for backward compatibility.

Q: I have just upgraded my workstation to HP-UX 10.01. I have noticed a kernel parameter that I do not recognize. The parameter is *no_lvm_disks*. What does it represent?

A: This parameter refers to the existence of *lvm* disks on the system. By default at boot time, the system checks for LVM data structures on the configured root, swap, and dump disks. If no LVM disks actually exist on your system, setting *no_lvm_disks* to 1 will speed up the boot process. The check for LVM structures will be skipped. If you are using LVM, then leave the parameter at its default value of 0.

Q: Can I move a file system disk from my 710 running 9.05 to my 735 running 10.10?

A: Yes, you should be able to utilize the disk on the 10.10 system. Simply mount the file system as an HFS file system.

Q: I'm preparing to upgrade my workstation from 9.05 to 10.01. I have already loaded the conversion tools on the system. What is the utility called *snoop*?

A: The *snoop* utility is used to check for conditions that would cause problems for your upgrade. The utility writes to a log file called */var/adm/sw/snoop.log*. You can actually look at this file while *snoop* is running. Each time you run *snoop*, it adds to the file rather than overwriting it. It also runs the upgrade in a preview mode. As a result, also check the log file */var/adm/sw/swagent.log* for errors and warnings.

The *snoop* utility divides its output into PROBLEMS and CAUTIONS. The PROBLEM issues must be acted upon for a successful upgrade. The CAUTION issues should be addressed, although they are less critical.

The utility will check such items as memory requirements, disk space, and cluster issues, to name a few. For more information about running *snoop*, refer to the manual *Upgrading from HP-UX 9.x to 10.x*, P-N B3782-90073.

Q: I have a mixed cluster environment that includes a 700 server, several 700 clients, and a few 300/400 clients. I'd like to upgrade my systems to 10.10. How can I continue to utilize all of my systems?

A: Since 10.x does not support Series 300 and 400 systems, you must remove any such clients from the DUX cluster before you upgrade the server to any 10.x release. Some options to

consider for those clients would be:

- convert the clients to stand-alone machines
- convert one client to a 300/400 cluster server
- upgrade the clients to Series 700 systems
- convert the clients to X-terminals

Converting to stand-alone is an option if the clients in question have at least one disk drive each. You can nfs mount non-system files from the 700 server. Converting to stand-alone necessitates a reinstall.

A Series 300/400 cluster server needs at least 16 MB of memory and at least a 420-MB disk for the OS and additional disk space for swap. You'll need about 30 MB per client unless it will swap locally. Again, you can nfs mount non-system files from the 700 server.

It is possible to upgrade some Series 400 systems to Series 700 machines. Check with your local HP office about the possibility of a board upgrade for your machine. If a processor upgrade is performed, the system should be added back to the diskless environment after the upgrade to 10.x.

To convert your Series 300 and 400 clients to X-terminals, you will need the XTERM300 product, which is part of the HP-UX 9.10 release. For more information, refer to the section "Using the Series 300/400 Xterminal Product" from the manual *HP-UX 9.x/10.x Interoperability Guide*, P-N 5963-8920.

Q: I am familiar with the autologout features of *csch(1)* and *ksh(1)*. Does HP VUE provide a similar feature?

A: This functionality is not "officially" offered in HP VUE. It is possible to implement this feature using the following steps:

1. Go into the VUE Style Manager (*vuestyle*)

- enable screen saver
- set the screen saver timeout value
- configure *vuesession* not to ask at logout
- disable *vuesession* confirmation

2. Create an executable script file called *~/vue/screensaver* with the following contents:

vueaction EXIT_SESSION

When VUE attempts to invoke the screen saver, the session will be terminated.

Q: How many sockets can a single process have open?

A: A socket open takes some of the same resources as a standard file open. So, the number of files and sockets that a process can have open is limited by the kernel parameter *maxfiles*. Notice that *maxfiles* is a soft limit. A process can increase its open file limit with the use of the *setrlimit(2)* system call. This increase is limited by the kernel parameter *maxfiles_lim*, which represents the per process hard limit.

Of course, the number of sockets and files will also be controlled by the kernel parameter *nfile*, which represents a systemwide limit.

Q: I have recently installed SoftBench on my 735 running HP-UX 9.05. I see that there are compilers under the SoftBench directory as well as in the system locations. What's the difference?

A: The compiler binaries you see under the SoftBench directory (*/opt/softbench* at 10.x and */usr/softbench* at 9.x) are actually wrappers for the standard compilers on your system. The SoftBench "compilers" are front-end interfaces that invoke the standard compilers with particular compile options. If you were to use those same options and directly compile your code with the standard compilers, the same object code would be produced.

Q: I have a 712/60 running HP-UX 9.05. This machine is the mail server for my network. The mail application uses an HFS file system to store mail messages. Occasionally, I have seen messages reporting that the file system is out of inodes. I have also seen messages that the inode table is full. What's the difference?

A: Let's first discuss the "out of inodes" message. The message will be displayed as:

`/<mount_point>: out of inodes`

When a file system is created, a static amount of file system space is set aside for inodes. Every directory and file in the file system is associated with an inode. The inode contains information about the file or directory such as mode, type, number of links, owner, size, etc. These inodes are often referred to as "disk inodes." The above message means that no more inode structures are available. The number of inodes created is based on the value of certain parameters when creating the file system, namely *nbpi*. See the *mkfs(1M)* man page for more details. When this limit is reached, the only options available are to remove files or re-create the file system and restore data.

With this in mind, it is very important to consider these parameters before the file system is created. If the file system will contain a small number of very large files, there's no reason to devote file system overhead to extra inodes. On the other hand, if there will be a large number of very small files, more than the default number of inodes may be needed. It is very possible to have plenty of disk space but no inodes to utilize it. There is a *-i* option with the *bdf(1M)* command to check inode utilization.

`inode: table is full`

This message refers to a kernel table that holds the "in-core inodes." When a file is accessed, the information from its disk inode is read into this table in memory. This "table" is actually a cache in that entries are not freed until needed. As a result, various utilities will often show that it is always full. When it is truly full, it will be impossible to access additional files.

The size of the table/cache is static and is defined by the kernel parameter *ninode*. To address this problem in the short term, you can eliminate extraneous processes that are accessing files. Long term, you may need to increase the value of *ninode*. This will not affect the system until a new kernel is generated and the system is booted off of that kernel. It is important to note that this modification will increase the size of the kernel. ■

General HP-UX questions are answered by Bill Hassell, a support engineer at the HP Atlanta Response Center. He can be contacted via e-mail at blh@hpuaatl.attl.hp.com. Workstation questions are answered by Susan Potter, an HP-UX system support engineer in the Atlanta Response Center. Her e-mail address is sup@atl.hp.com.

HP-UX System Administration

by Chris Curtin

Installing Software on HP-UX 10.x

WELL, I FINALLY INSTALLED 10.10 on one of my systems. (Aren't customers who DON'T want to upgrade great?) One of the first things I noticed during the install was the software installation tool *swinstall*.

swinstall is one of a set of applications for installing, removing, configuring, and monitoring the installed software on your systems. Called SD-UX (Software Distributor For HP-UX), it takes the place of */etc/updist*, */etc/update*, and */etc/netdistd*.

What's a Depot?

The *swinstall* program is used to install applications from a depot. A depot is a physical area, either on disk or CD-ROM or tape, that contains the software applications available for installation. Unlike 8.x and 9.x, the *swinstall* program handles all three kinds of depots. (Remember in 9.x using */etc/updist*, */etc/update*, and */etc/patch_install*? The functionality of all three is handled by *swinstall*.)

The disk depot, also known as a directory depot, is where applications you need to distribute via the network can be loaded once so that the networked systems can access them. The default directory for this is */var/spool/sw*. (Again, think functionally of */etc/updist* and */netdist*.)

A CD-ROM drive is also considered a directory depot, though I would expect copying the applications from CD-ROM to the local disk would make for faster network updates. I have no hard data to back this up, but CD-ROM access speeds are nowhere near the SCSI speeds.

The tape depot uses a tape drive as its location. Unlike directory depots, tape depots cannot be used to update network systems. You must first copy the applications to a directory depot. For a local installation, I skipped copying the

tape to the directory depot.

A quick tangent: when you are installing HP-UX 10.x, *do not* use the default sizes for */var* or */opt* if you plan on doing anything useful. Unlike 9.x, which installed most applications into the */usr* directory, 10.x split applications between these two. I was burned by a too-small */opt* directory when installing all my compilers (*swinstall* noticed there wasn't enough disk space and didn't try to install, but I still had to reconfigure my disks.)

Registering Depots

Unlike 9.x, *swinstall* lets you have multiple depots configured at the same time. But, also unlike 9.x, you cannot mix Series 700 and 800 applications in the same depot.

If you want to set up a default depot for installation of software, you can either let *swcopy* do it (more about *swcopy* later) or you can do it via the command line. I recommend doing it via the command line since if you don't exactly remember the path to use, *swcopy* can be hard to use.

Registering a depot via the command line uses the *swreg* command. For example, to register your CD-ROM as a depot (which by default it is not, even though HP is pushing people towards CD-ROM-only distribution) use:

```
/usr/sbin/swreg -l depot @ /cdrom
```

assuming */cdrom* is your mount point for your CD-ROM. (Note: In 10.x HP recommends mounting any non-local disk under */mnt*. I, however, don't think of the CD-ROM as non-local. It is a religious issue.)

The *-l* depot option says that we are registering a depot; there are some other

OUR 16th YEAR

QUALITY HP WORKSTATIONS - WITHOUT THE HP PRICE

RENT · LEASE · BUY

Immediate Delivery On Most Items

- 700 SERIES: Models 710, 715/50, 715/75, 715/100, 735, 735/125, 712/60, 712/80, 745i
- X-STATIONS: ENVIZEX, C270X
- All 300, 400T and 400E Series
- Memory, Features & Disc Upgrades for all Workstations
- CPU Upgrades: 715/50, 715/75, 735, 735/125, 425, 380, 360

*More Than 1500 Satisfied Customers
Throughout The USA and Worldwide*

For technical info, specs or pricing
call Mordy or Carol



C.S.U. Industries, Inc.

207 Rockaway Turnpike, Lawrence, NY 11559
(516) 239-4310 FAX (516) 239-8374

options, but they are pretty advanced and you probably won't need them. The @ is required and the last parameter is the directory to register as a depot.

Unregistering a depot is as easy as adding a -u option to the same command line.

Codewords

We all remember the codeword policy in 9.x: The applications linked to the hardware ID. If you had a CD-ROM, the codeword also included the disk part number in its key. I've had serious problems with this scheme—like losing a server, getting a new CPU card from HP and having the SE forget to reset the hardware ID so that I couldn't load anything from the install media!

In 10.x HP is taking an honor approach. You still get a codeword for the products you have purchased. You still need the original CD-ROMs that match the codeword, but instead of the hardware ID, the other part of the key is your customer support identifier. That's right, no hardware ID is required.

I really hope you didn't just think, Cool, now I can install that compiler on all my workstations. As a software developer, I can assure you that software piracy is a big issue that costs this industry billions of dollars a year, so I discourage that kind of thinking (and acting).

As HP puts in the manuals, "It is your responsibility to ensure that the codeword and software are used in this manner" (page 1-9 of the *Managing HP-UX Software With SD-UX* manual). So don't load software into a depot and set it on the Internet with a README file containing the support ID and codewords.

Structure of An Application

In 9.x, using *fpkg* to create the update

CIRCLE 140 ON READER SERVICE CARD

tapes, you really had only two levels of groupings: applications and filesets. Usually you had several dozen top-level applications in one product. For example the C++ compiler included man pages, a debugger, a linker, the compiler, and the contributed libraries. If you didn't know what you were doing, or were not careful, you could install too much stuff or miss a critical fileset. (How many of your systems have the Japanese language filesets installed? Assuming you're not Japanese of course.)

SD-UX has four levels: Bundle, Products, Subproducts, and Filesets. A Bundle is an HP-only item. You cannot use the development tools in the SD-UX toolset to create them. I look at a bundle as a grouping together of products that HP thinks are important. For example, on your installation CD-ROM, HP has Bundles for the various languages. So instead of hunting through all the filesets to remove those languages not required (or to install some combination) HP lets you select a Bundle and your done.

After the Bundle are the Products. The Product is the top level of a customer SD-UX distribution and typically where HP puts the products that you have purchased from them. (Some exceptions are options to the compilers like Blink Link or some of the other development environments. These are in subproducts or filesets.)

Products do not have to be for only one revision of the software or operating system. A Product can have several Subproducts for a specific release of the Product or for a specific hardware platform.

Pretty neat if you ship software to multiple customers on multiple OS releases on different platforms. Just cut

the full tape and let *swinstall* and the customers tell you what they want. *swinstall* does detect the hardware it is on and show you only the applications for the specific hardware.

Subproducts can also be used to group multiple Filesets together. A Fileset is the actual applications or components to install. This level is very similar to the old */etc/update* layout.

SD-UX Applications

There are four primary applications in SD-UX that you will need to use under normal circumstances: *swinstall*, *swcopy*, *swremove*, and *swlist*. All four are installed under */usr/sbin/*. (*/usr/sbin/* is where system binary or executables are placed in 10.x.)

swinstall is used for installing software from a depot to a local machine. The depot can be on the same machine or on a networked server.

swcopy is used for copying software from CD-ROMs or tapes to a depot. The depot must reside on the same machine as the tape or CD-ROM drive.

swremove removes software from the depot. The depot must reside on the same machine the command is executed on.

swlist lists the installed software on your machine. No more *grep*ing through */system* or */etc/* filesets to see what you have installed!

swinstall, *swcopy*, and *swremove* have both a text interface and an X Window System interface. They behave similar to SAM, *update*, and *updist* in 9.x and the interface is identical to that of SAM in 10.x.

There are several other applications, including *swconfig* (use this to remove applications) in the SD-UX toolkit, but they are fairly advanced and I'll try to cover them in the coming months.

One warning though: using the enter

or return key to select an object in *swinstall* or the others also opens the object. I kept using the enter key (instead of the space bar) to mark and unmark objects and I would find myself in the next level down or in some copyright/fileset description screen.

HP-UX 10.10 First Impressions

Now that I have a 10.10 system, I have several first impressions:

1. Finally a "normal" UNIX file system layout!
2. The *HP-UX System Administration Tasks* manual is about a third the size of the 9.x version. I don't know if they have moved topics to other manuals, but I was surprised by this.
3. There is a host of interesting applications and features such as a network time server, a supported automounter, a new version of SAM, and a number of other applications started at boot time that I haven't had time to investigate yet. Watch your system boot and see everything that starts or attempts to start!

That's it for this time. Since I am new to 10.x, please send your comments and suggestions for future columns and tips to me via e-mail. As always, keep the e-mail coming. ■

Chris Curtin, a software developer for Bradley Ward Systems, Inc. in Atlanta, Georgia, specializes in device driver development for factory automation on the HP 9000. He can be reached via e-mail at: chris@bwilab3.atl.ga.us.

Stuck using tar?

(or fbackup, cpio, or dump?)

What a sticky mess! Standard UNIX backup utilities force you to glue on scripts to make them work right, have performance like molasses, user interfaces that are clear as pitch, and reliability that could drag your career down into a pit. It's time to kick asphalt. Load BACKUP/9000 (it takes less than 10 minutes), and automate super fast, reliable backups and restores on any networked system via its slick user interface. Let BACKUP/9000 back up your Oracle databases hot, track tapes and files, manage media, schedule backups, etc. Don't get stuck - get something faster, easier, and more reliable. Contact us for a free demo today, before things get really messy.



+1-800-89ORBIT
+1-510-837-4143
Fax +1-510-837-5752
info@orbisw.com
www.orbisw.com



CIRCLE 72 ON READER SERVICE CARD



by David L. Totsch

Fast Symlinks

MANY CONVERSATIONS TODAY EVENTUALLY find their way to the topic of performance. I am not yet sure if we just like to talk about performance so that we can brag, but many are seeking ways to squeeze some extra performance out of their systems. Since few of us have obvious bottleneck battles to wage, we tend to discuss the small skirmishes we have engaged in. One of those small skirmishes has to do with large numbers of frequently used symbolic links.

Normally, a symbolic link uses one disk block to store the path name to the destination. Under HP-UX 10.x it is possible to instruct the system to store the path name within the inode structure where the inode would normally store direct addresses for disk blocks. This is all given that the path name is short enough to fit within those fifteen address slots. If it fits, the system will not have to incur the disk overhead of reading that one block to obtain the path. If the path does not fit, it uses the standard method of allocating a disk block to store the path. Since there are 15 block addresses, fast symbolic links will work for link paths of 60 characters or less (the fifteen blocks hold 32-bit address—each address can hold 4 characters).

Here is a listing of a directory:

```
9987 lrwxrwxrwx 1 dlt users 66 Oct 21 16:06 CP ->
    ../../../../usr/vue/config/types/tools/System_Admin/ChangePassword
9985 -rw-rw-rw- 1 dlt users 39 Oct 21 14:31 myFile
9986 lrwxrwxrwx 1 dlt users 6 Oct 21 09:30 myLink -> myFile
```

To prove that the information is stored in the inode, here is the output of *fsdb(1m)* (remember that the angels leap off of your shoulders when you use *fsdb*):

```
9986i
i#:9986 md: l--rwxrwxrwx ln: 1 uid: 329 gid: 100 sz: 6 ci:0
symlink: myFile

at: Mon Oct 21 14:12:31 1996
mt: Mon Oct 21 09:30:02 1996
ct: Mon Oct 21 09:30:02 1996
9987i
i#:9987 md: l--rwxrwxrwx ln: 1 uid: 329 gid: 100 sz: 66 ci:0
a0 : 65322 a1 : 0 a2 : 0 a3 : 0 a4 : 0 a5 : 0
a6 : 0 a7 : 0 a8 : 0 a9 : 0 a10: 0 a11: 0
a12: 0 a13: 0 a14: 0
at: Mon Oct 21 16:06:11 1996
mt: Mon Oct 21 16:06:07 1996
ct: Mon Oct 21 16:06:07 1996
```

From the above, the file *myLink* is a symbolic link to *myFile*.

Since the information is stored within the address, a0-14 are not displayed and the symlink line is printed instead. Although the file CP is also a symbolic link, it exceeds the 60-character limit and the path is stored on a separate disk block.

Total Internet

- Statistical Reporting
- Site Hosting
- On-Line Ordering
- Fax-Back Info

Faxing Solutions

- Host-Fax - HP 3000
- Replix - HP 9000/UNIX
- LanFax - PC Networks


DataComm Consolidation

- Bridges, Mux, CSU/DSU
in one expandable box

High Availability

- Disk Mirroring
- Disk Shadowing

I/O Data Systems, Inc.



HP 3000 HARDWARE

HP 9000 HARDWARE

HP CERTIFIED



I/O DATA SYSTEMS, INC.

Ph. (216) 835-2211 Fax (216) 835-0220

"An Inc. 500 Company"

CIRCLE 171 ON READER SERVICE CARD

That extra disk access may not sound like much overhead. But, if you have more than a handful of symbolic links on your system that are heavily used by an application, the savings may be of interest. More directly, if you have an HP-UX 10.x system and have installed the symbolic links to support a legacy application, you could see a modest improvement in performance. There is a caveat, though. You must set the kernel parameter `create_fastlinks` to 1 and remake your kernel.

Newly created symbolic links will be fast links if the path is less than 60 characters. If you have already installed the transition links on your HP-UX 10.x system, you will need to remove and reinstall them. ■

Large Databases Require New Solutions for Backup !!!!!

FINALLY! Backup 4 to 8 Gigabytes/Hour/Drive
Store 20 to 40 Gigabytes/Tape
Random Libraries from 5 to 264 Cartridges


DallasTools

Back-up, Restore and Archiving
Software for Heterogeneous Networks

Dallastone also sells and supports all major brands –
Quantum, Breece Hill, ADIC and Odetics



DALLASTONE
2 Cote Lane
Bedford, NH 03110



Phone 603-647-8168
Fax 603-624-2466
Email dttool@delphi.com

After serving several different organizations over the past seven years as a system administrator with various flavors of UNIX, David L. Totsch still enjoys the profession. He also enjoys discussing UNIX with just about anyone. At present, he is working with HP-UX systems and wide-area networks for a Fortune 100 company in the Piedmont area of North Carolina. He can be reached via Internet: dtotsch@wfu.edu.

CIRCLE 76 ON READER SERVICE CARD



PVM—PARALLEL COMPUTING ON YOUR DESKTOP

by Weicheng Jiang

Parallel computing is a cost-effective way to gain more processing power and memory than available on a single workstation. Multiprocessor workstations and servers are now commonplace, and high-speed networking technologies such as Fibre-Channel allow you to build your own parallel computers.

Cluster computing also has the added benefit of fault-tolerance. If one machine goes down, the rest of them can continue to run without interruption. Developing applications to run in parallel is no easy task, however. Parallel Virtual Machine (PVM) is a programming environment designed to help you write portable parallel applications.

What Is PVM

PVM is software that enables execution of parallel applications across multiple nodes on a network. With this software a user can turn a loose group of machines into a parallel computer. PVM runs on most UNIX machines, and on any network that supports the TCP/IP protocol. PVM can be started from any one machine. If the user supplies a host file with a list of machine names, those machines will be added to the Virtual Machine (VM) configuration at start-up. A virtual console can be brought up on any host in the configuration to monitor the status of the VM. The user can issue commands from the console to add or delete hosts from the VM, and to list active jobs. New jobs can also be spawned from the console.

For parallel applications, PVM has a simple message-passing interface for exchanging data between different tasks. Each task is identified by a unique global task ID. The IDs of the sender and recipient are encoded in the message header and the message is routed to the appropriate task by the PVM daemons on the source and destination hosts.

PVM can also deliver UNIX signals to a task, regardless of its physical location.

The Power of PVM

PVM is a low-cost, versatile tool for parallel computing. At one extreme, it can turn a loosely scattered, under-utilized set of UNIX workstations into a powerful parallel computer. At the other extreme, the user can build a supercomputer from stripped-down UNIX boxes, connecting them through a fiber-optic network.

PVM is also supported on massively parallel computers such as the Cray T3D, IBM SP-2, and Intel Paragon. The

most important advantage of PVM is its portability. A program written in PVM can run on almost any hardware in use today, from PCs to supercomputers. This removes hardware dependency from an application and reduces the cost for development and future upgrades. An application can also be developed on a desktop system and then moved to a supercomputer for production runs.

PVM also works in a heterogeneous environment that consists of systems from different vendors. The programmer does not have to worry about incompatible data formats when passing data between machines with different architectures; PVM will do the necessary conversion. Workstations can be linked up with multiprocessor computers to form a giant virtual machine. Results generated by a supercomputer can also be displayed on a graphics workstation via PVM.

For users new to parallel computing, PVM has a small set of functions that are intuitive and easy to use. The concept of a virtual machine also makes programming simple.

PVM has been available for several years. It has gained wide acceptance among technical users, and has become a de facto standard. Computer vendors have recognized this trend and started to offer PVM on their systems.

A rich set of tools is available to PVM users. Visualization tools such as XPVM and VPE provide a graphical interface to PVM and allow the user to monitor the performance of a PVM program. From a replay of the trace data recorded during an execution of the program, the user can identify any performance bottlenecks and fine-tune the code. Parallel debuggers for PVM are also available from software vendors.

How PVM Works

PVM consists of a daemon and a message-passing library (*Figure 1*). A daemon is started on every host in the Virtual Machine. Users' programs need to be linked with the PVM library at compile time. There are three ways to start a new PVM task:

- run it like any other UNIX process
- spawn it from the console
- spawn it from another PVM task.

Normally the PVM daemon forks and execs the new process. The process then enrolls in PVM and gets an ID from the daemon. A TCP socket connection is established between the task and the daemon. (On a shared-memory multiprocessor system, a shared buffer is used to pass messages.) The new task can query the daemon for information on other tasks and the configuration of the Virtual Machine.

When a task sends a message to another task, the message is usually routed by the local daemon. The daemon decodes the message header and forwards the message to the destination host. The daemon on that host then passes the message along to the intended recipient. Daemons on different hosts communicate with each other via UDP sockets. They will resend any packets lost by UDP.

The programming model of PVM is quite simple. Each task is identified by a unique ID. From the programmer's point of view, it really doesn't matter where the task is running. This model is similar to multitasking on a mainframe. The PVM console gives the user a global view of the Virtual Machine; commands can be issued there to query the status of any task or to send a signal to a particular task.

Continued on Page 33

FIGURE 1 *PVM Implementation on EPS Fibre-Channel Cluster*





FIGURE 2 *PVM Intra-host and Inter-host Communications*



The most important concept in PVM is the single-system image. It hides the hardware details from the programmers, and enables them to view the distributed environment as a monolithic system.

The sample program shown in *Listing 1* illustrates the basics of PVM. A sample run of this program from the PVM console is shown in *Listing 2*. We list all commands at the start to help the reader understand what we are doing. PVM was first started on *April*, and two other hosts, *Motel6* and *Badger*, were added from the console. Three instances of *greet* were then spawned, and the outputs were forwarded back to the console.

HP-PVM

In response to customer demand, HP now supports PVM across its UNIX product line, from the EPS Fibre-Channel Clusters to the Exemplar SPP supercomputers. Our implementation of PVM offers low latency and high data throughput. *Table 1* shows how PVM performs on an EPS Fibre-Channel Cluster. (Units are microseconds and megabytes per second, respectively.) The bandwidth is for 80-Kbyte messages. Tasks on the same host communicate using shared memory, while off-host messages are routed through the Fibre-Channel connections.

To achieve good performance, messages are sent directly to recipients, bypassing the PVM daemons. The daemon is still needed to help set up the communication channel. Once that's done, the daemon will go to sleep. *Figure 2* illustrates the message routes.

The numbers in the table show that shared memory is much faster than the network. We expect this imbalance to be corrected when the Gigabit Fibre-Channel becomes available. But even a

LISTING 1 Sample PVM Program

The simple program below illustrates the basics of PVM:

```
#include <stdio.h>
#include <string.h>
#include <strings.h>
#include "pvm3.h"

#define TAG      1  /* message tag */
int mytid;         /* my task ID */

main()
{
    int ntask;                /* number of tasks */
    struct pvmtaskinfo *tasks; /* information about tasks */
    char msg[128];            /* greeting message */
    int i;

    if ((mytid = pvm_mytid()) < 0)
        exit(-1);            /* enrollment failed */

    sleep(1);                 /* wait for everyone */
    pvm_tasks(0, &ntask, &tasks); /* find out who's here */

    sprintf(msg, "Greetings from %x on ", mytid);
    gethostname(msg + strlen(msg), 64);

    for (i = 0; i < ntask; i++)
        if (tasks[i].ti_tid != mytid && tasks[i].ti_a_out[0])
            /*
             * Send a nice greeting to everyone
             */
            pvm_psend(tasks[i].ti_tid, TAG, msg, strlen(msg), PVM_BYTE);

    for (i = 2; i < ntask; i++) {
        /*
         * Collect greetings from everyone
         */
        bzero(msg, sizeof(msg));
        pvm_precv(-1, TAG, msg, sizeof(msg), PVM_BYTE, 0, 0, 0);
        puts(msg);
    }

    pvm_exit();
    exit(0);
}
```


LISTING 2 *Sample Run*

```

april>pvm
pvm> help
help - Print helpful information about a command
Syntax: help [ command ]
Commands are:
  add - Add hosts to virtual machine
  alias - Define/list command aliases
  conf - List virtual machine configuration
  delete - Delete hosts from virtual machine
  echo - Echo arguments
  halt - Stop pvmds
  help - Print helpful information about a command
  id - Print console task id
  jobs - Display list of running jobs
  kill - Terminate tasks
  mstat - Show status of hosts
  ps - List tasks
  pstat - Show status of tasks
  quit - Exit console
  reset - Kill all tasks
  setenv - Display or set environment variables
  sig - Send signal to task
  spawn - Spawn task
  trace - Set/display trace event mask
  unalias - Undefine command alias
  version - Show libpvm version
pvm> add motel6
1 successful
      HOST      DTID
motel6      80000
pvm> add badger
1 successful
      HOST      DTID
badger      c0000
pvm> conf

```

LISTING 2 *Sample Run, continued*

```

3 hosts, 1 data format
      HOST      DTID      ARCH      SPEED
      april     40000     HPPA      1000
      motel6    80000     HPPA      1000
      badger    c0000     HPPA      1000
pvm> spawn -3 -> greet
[1]
3 successful
t80001
tc0001
t40002
pvm> ps -ax
      HOST      TID      FLAG 0x COMMAND
      april     (cons)      4/c -
      april     40002      6/c,f greet
      motel6    80001      6/c,f greet
      badger    c0001      6/c,f greet
pvm>
[1:t40002] Greetings from c0001 on badger
[1:t40002] Greetings from 80001 on motel6
[1:t80001] Greetings from 40002 on april
[1:t80001] Greetings from c0001 on badger
[1:t40002] EOF
[1:tc0001] Greetings from 40002 on april
[1:tc0001] Greetings from 80001 on motel6
[1:t80001] EOF
[1:tc0001] EOF
[1] finished
pvm> ps -ax
      HOST      TID      FLAG 0x COMMAND
      april     (cons)      4/c -
pvm> halt
april>

```

bandwidth of 20 MB/s is sufficient for most parallel applications, considering the CPU is running at only 180 MHz. (A CPU on average executes one instruction per cycle.) Additionally, the shared-memory bus is used for memory access and I/O by all the processors, while the Fibre-Channel is dedicated to inter-host communications. For large messages (1 MB or more), the performance of shared memory falls off rapidly, while the throughput of Fibre-

Channel improves slightly.

To see how PVM performs in a real-world program, we ran a 3-D multigrid parallel benchmark, which is described as "very communication intensive," from NASA Ames Research Center. The results are listed in *Table 2*. The benchmark was run on an EPS Fibre-Channel Cluster with two K400s, each with four 100-MHz PA-RISC CPUs. For comparison, we also ran the same benchmark on four 712 workstations, connected by

Ethernet, each with an 80-MHz CPU. The numbers are in the last column.

The high communication overhead on the 712s, which accounted for 63 percent of the total runtime, is due entirely to the low data rate of the Ethernet. In contrast, the communication overhead on the EPS cluster is less than 20 percent, and the speedup scales well with the number of processors. This proves the EPS cluster, combined with PVM, is an effective platform for parallel computing.

The problem size we chose for the benchmark was 128³. We couldn't run it at the maximum size of 256³ on one K400 with 256 Megabyte of main memory, but we were able to run it at that size on the cluster. This illustrates another advantage of parallel computing: it gives you access to more memory. For large programs, memory is as much a concern as is CPU speed.

Summary and References

While a parallel program can be written with only a few of the functions in the PVM library, PVM offers many more features for advanced users. For example, there are functions that enable fault tolerance and load balancing to be built into a user's program. There is also a group library for collective communications.

Third-party tools for performing checkpointing, process migration, or

scheduling in PVM programs are available from various research organizations


and software vendors.

To learn more about PVM, please visit the PVM home page at the Oak Ridge National Lab web site:

<http://www.epm.ornl.gov/pvm>

If you don't have World Wide Web access, a book on PVM is also available: *PVM—A User's Guide and Tutorial for Networked Parallel Computing*, Al Geist, et al., MIT Press, 1994. □

Dr. Weicheng Jiang received his Ph.D. in Applied Physics from Dartmouth College in 1992. He worked on the PVM Team at the University of Tennessee until Oct. 1995, when he joined the Chelmsford System Software Lab of Hewlett-Packard Co. He has implemented several optimized versions of PVM for various multiprocessor systems, including the Intel Paragon, IBM SP2, and Cray Superserver. He is also a coauthor of the PVM book from MIT Press: PVM—Parallel Virtual Machine.



HP/APOLLO[®] Sun Microsystems[®] WORKSTATIONS

■ Quality depot repairs at substantially reduced cost

**On-Site Maintenance
for
HP/Apollo - Sun Microsystems
Savings up to 50%**

**Systems/Spare Part Sales
Technical Assistance
Buy ■ Sell ■ Repair ■ Service
Satisfaction Guaranteed**

CALL US TODAY AT AMC COMPUTER SERVICES, INC.

146-B Rangeway Rd., P.O. Box 286

N. Billerica, MA 01862

508-670-9395 ■ FAX 508-670-9327

1-800-878-2627

E-MAIL: AMC@TIAC.NET

CIRCLE 3 ON READER SERVICE CARD

TABLE 1 Performance of PVM on Fibre-Channel Cluster

	SHARED-MEMORY	FIBRE-CHANNEL
Latency (μs)	7	152
Bandwidth (MB/s)	154	20

TABLE 2 Performance of NAS Multigrid Benchmark in PVM

	K-400 CLUSTER				712s
Number of Processors	1	2	4	8	4
Run Time (seconds)	46.4	23.2	13.1	7.2	41.7
Communication Time (seconds)	0	0.43	0.8	1.4	26.2

10.X

What's new with

Startup

by John Fenwick

and Configuration

Introduction

This article explains the 10.X system startup and configuration models and outlines the differences between HP-UX Releases 9.X and 10.X. All of the changes apply to both Series 700 and Series 800 computer systems. These differences will primarily affect software developers and system administrators. Software developers will construct rc files in a different manner, specifying run-level execution order and enabling control through configuration variables. System administrators will control subsystem behavior, on a per-host basis, by modifying variables that control subsystem startup and shutdown.

The 10.X Startup and Configuration models introduce entirely new designs for subsystem control. Users no longer modify large `/etc/rc*` scripts to modify the system behavior. Instead, rc files have been separated into functional pieces. Developers and users need only modify configuration files to control the system startup and shutdown behavior. If your

system or application modified or created its own `/etc/rc*` script, this will need to be modified for 10.X. The startup and configuration specifications apply to all system components, including both the operating system and installed applications.

Design Considerations for 10.X Startup and Configuration Models

The 10.X Startup and Configuration models are similar to the designs used by other major UNIX vendors. Important new features of the design include:

- The model separates the execution scripts from the configuration information required for execution. System administrators may easily modify the behavior of the startup/shutdown sequence by changing configuration variables.
- As a result of the separation of Operating System code from user-entered configuration values, the design solves

the problem of correctly updating modified system configuration files.

- Individual subsystems may now be started or stopped on a per-runstate basis. This enables fine levels of control and subsystem separation.
- All subsystems are now treated uniformly at startup. There are well-defined interfaces for input parameters and exit values for startup scripts. This makes the design easy to extend and modify.

HP-UX Bootstrap Strategy

HP-UX uses a “bootstrap” approach to startup. In this paradigm, the operating system comes up through a series of steps of increasing functionality, resulting in a fully functioning, network-capable system.

In the early stages of Kernel and User Space boot, all required files for the kernel, necessary commands, and required configuration files must be located on the root disk. In the 10.X file system layout, the locations of Required Root Disk files are shown in *Figure 1*. The Subsystem Startup Files used for bringing up individual subsystems (the subject of this article) are all found within the */sbin* and */etc* directories as noted in the diagram.

Startup/Shutdown Overview

The old */etc/rc*-based startup and configuration design has been replaced in the HP-UX 10.X release. The new startup/shutdown design consists of three parts:

1. Execution Scripts used to start up and shut down individual subsystems are contained in */sbin/init.d*.
2. Configuration Files for each execution script are contained in */etc/rc.config.d* and reside in a filename that corresponds to the execution script which they configure.
3. Link Files in */sbin/rc*.d* control the sequencing order of the execution scripts. The ordering of these link files should not be modified by the user.

Figure 2, “9.X - 10.X RC Configuration Mapping,” shows the relationship of 9.X files to 10.X files. Prior to 10.X, both script code and configuration information, such as IP addresses and hostnames, were contained within the */etc/*rc** files. The new model partitions the configuration data from the scripts. Administrators use the configuration variables in */etc/rc.config.d* to change the behavior of scripts in */sbin/init.d*. Script sequenc-

ing is also a new feature for 10.X and is controlled through link files in */sbin/rc*.d* directories.

Execution Scripts

The */sbin/init.d* directory contains all scripts used to start up and shut down various subsystems. No script may invoke any of the other scripts in this directory. Scripts obtain configuration data from variables in */etc/rc.config.d* (more on this later), which must be sourced by the execution script. All files in the */etc/rc.config.d* directory may be read by sourcing the single file */etc/rc.config*.

The startup and shutdown scripts installed with the operating system in */sbin/init.d* are not to be edited. Any changes made to these scripts will be overwritten when a new software release is installed. Modifying the behavior of a subsystem script is accomplished by using configuration variables, discussed in “Configuration Variable Scripts.”

In general, each script under */sbin/init.d* should perform both the startup and shutdown functions. In order to control the functionality within the script, each must also support standard arguments and exit codes. Scripts must be written for the POSIX shell. A template script that can be copied and modified for use in a subsystem may be found in */sbin/init.d/template*.

Arguments to Scripts

The startup/shutdown scripts must recognize the following four arguments:

- *start_msg*. The *start_msg* argument is passed to scripts so the scripts can report back a short message indicating what the “start” action will do. For example, when the lp spooler script is invoked with a *start_msg* argument, it echoes “Start print spooler.” This string is used in the startup checklist. Note that when given just the *start_msg* argument, scripts will only echo a message and NOT perform any other actions.
- *start*. Upon receiving the *start* argument, the script should start the subsystem. All output should be echoed to *stdout*.
- *stop_msg*. The *stop_msg* argument is passed to scripts so that they can report back a short message indicating what the “stop” action will do. For instance, when the lp spooler script is invoked with a *stop_msg* argument, it echoes “Stop print spooler.” This string is used in the shutdown checklist. Note that when given just the *stop_msg* argument, scripts will only echo a message and NOT perform any other actions.

Continued on Page 40

FIGURE 1 Locations of Root Disk and Subsystem Startup Files

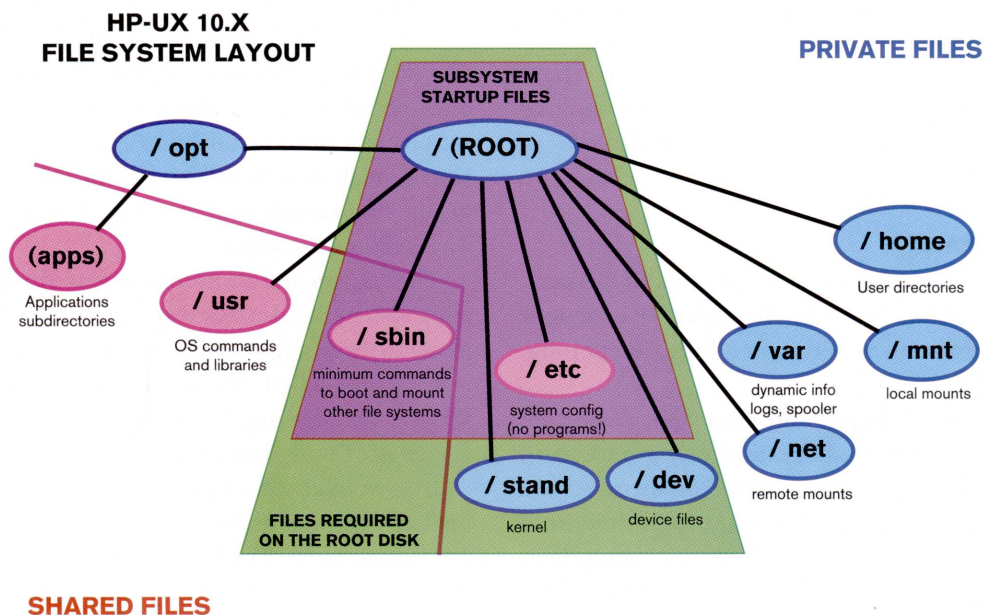
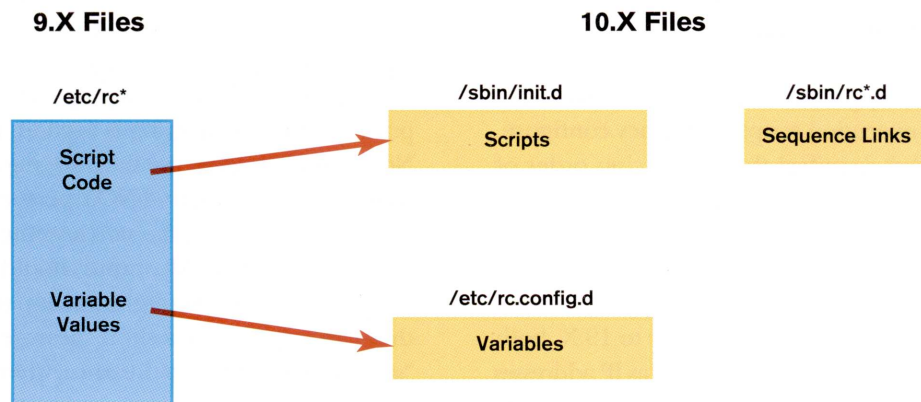



FIGURE 2 9.X - 10.X RC Configuration Mapping



When you run your Data Center right, you are a winner!



DESIGN
3000
PLUS
INCORPORATED

When your Data Center has problems it's a race against time until they are corrected. We have JMS to schedule and control your batch jobs, and CALLBACK to monitor and notify you of critical conditions. You will be notified by phone, pager or alpha pager of any user defined conditions or other problems that occur within your Data Center.

JMS JOB MANAGEMENT SYSTEM

JMS puts you in control of your Data Center with scheduling and execution of batch jobs. JMS has no command language to learn so it's easy to use and yet has sophisticated scheduling capabilities that are second to none. JMS network capabilities allow scheduling and job dependencies across the network. Restriction features allow you to limit those pesky system hogs. Ad hoc streamed jobs may be incorporated to allow complete batch job control.

CALLBACK SYSTEM DOWN DETECTOR AND JOB NOTIFICATION SYSTEM

Voice/Beeper/Digital or our new Alpha Pager notification for:

- Adverse Temperature • Power Loss • System Hangs • Job Aborts • Printer Status • Physical Conditions • User Requests • Logon Security • Reply Pending • Console Messages • Special Job Events • Spoolfile Scanning • And much more...

We are the leaders in tools for the "Lights Out Environment." Call us today for a free 30 day demo. MPE/iX or HP-UX.

Corporate Headquarters

3470 Pipebend Pl. NE • Suite 120 • P.O. Box 13086 • Salem, OR 97309-1086
Phone: (503) 585-0512 • FAX: (503) 585-1706 • E mail: Design3000@aol.com

International Sales Office

System Software Intl. • Oakmoore Court • Kingswood Road • Hampton Lovett
Droitwich, WR9 0QH, UK • Phone: +44 (0) 1905 794646 • FAX: +44 (0) 1905 794505



CIRCLE 44 ON READER SERVICE CARD

- *stop*. Upon receiving the *stop* argument, the script should shut down the subsystem. All output should be echoed to *stdout*.

When passed the start and stop arguments, an execution script should not echo any messages indicating the entry or exit from the script. *start_msg* and *stop_msg* arguments will be passed to the execution scripts by */sbin/rc* to record these messages both on the console and in log files, indicating the execution of the script.

Naming Conventions

The startup and shutdown scripts are named after the subsystem they control. For example, the */sbin/init.d/cron* script controls the *cron* daemon.

Scripts and Console Output

To ensure proper reporting of startup events, startup scripts will be required to comply with a few guidelines for script output and exit values.

The messages echoed by the execution script when *start_msg* and *stop_msg* arguments are passed should contain a single-line message with no more than 30 characters.

Status messages, such as "Starting Subsystem daemon," must be directed to *stdout*. All error messages must be directed to *stderr*. Both *stdout* and *stderr* are redirected to the log file */etc/rc.log*, unless the startup checklist mode is set to the raw mode. In this case, both *stdout* and *stderr* output go to the console.

Startup scripts, and the daemons or binaries they execute, must not send messages directly to the console during system boot or shutdown. This restriction exists because console output during the boot or shutdown sequence will overwrite the graphical checklist, leaving the checklist unreadable. These messages should be directed to *stdout* or *stderr*. The display mode is controlled by the assignment of the variable *LIST_MODE* in the file */etc/rc.config.d/list_mode*.

Exit Values

Exit values for startup scripts are as follows:

- 0—script exited without error. This causes the status "OK" to appear in the checklist.
- 1—script encountered errors. This causes the status "FAIL" to appear in the checklist.
- 2—script was skipped due to overriding control variables from */etc/rc.config.d* files or for other reasons, and did not actually do anything. This causes the status "N/A" to appear in the checklist.

3—script executed normally and requires an immediate system reboot for the changes to take effect. This is reserved for key system components.

These are the only acceptable exit values. Returning an arbitrary non-zero exit value from a command to indicate failure may cause the script to appear to have been skipped in the checklist, or may cause the system to reboot.

Configuration Variable Scripts

Instead of spreading configuration data throughout the various rc files in the system, configuration data is structured as a directory of files that allows developers to create and manage their own configuration files, without the complications of shared file ownership. The directory that holds the configuration variable scripts is */etc/rc.config.d*. The configuration variable scripts will be sourced by the startup and shutdown execution scripts in */sbin/init.d* during system startup and shutdown. This configuration information is used by the execution scripts to enable/disable and configure subsystems (such as IP addresses).

Examples of these variables include:

- *HOSTNAME*: Internet name of your system.
- *IP_ADDRESS[0]*: Internet address of your system, in dot format.

These configuration variables may be defined as either simple shell variables of type string, or as singly indexed array variables of base type string. The use of indexed variables adds a powerful, yet easy to understand and use, capability to the system.

Consider a system that contains two (or more) network interfaces and may serve as a routing agent. The various network interfaces can be simply defined and manipulated as follows:

```
INTERFACE_NAME[0]=lan0
IP_ADDRESS[0]=15.13.185.120
SUBNET_MASK[0]=255.255.248.0

LANCONFIG_ARGS[0]="ether ieee"
INTERFACE_NAME[1]=lan3
IP_ADDRESS[1]=15.27.233.2
SUBNET_MASK[1]=255.255.248.0
LANCONFIG_ARGS[1]="ether ieee"
```


Each execution script may reference its variables in one of two ways. If the execution script references only the variables delivered with its product or file set, it may explicitly source the file in */etc/rc.config.d/<subsystem>*. If the script requires variables that are delivered by other products or file sets, it may just source */etc/rc.config*, which is a script that sources all the files below */etc/rc.config.d*. The master configuration file */etc/rc.config* skips files named “core,” or containing the characters [*.*, *~*, *#*] when it sources the contents of */etc/rc.config.d/*, to avoid sourcing a binary core dump file or a backup file that may have been left in that directory by a system administrator.

There must be no requirements on the order of the files sourced. This means configuration files must not refer to variables defined in other configuration files, since there is no guarantee that the variable being referenced is defined. Also, no registry of variable names is maintained, so developers should take care to avoid variable name-space collisions.

Configuration variable scripts are written for the POSIX shell (*/usr/bin/sh* or */sbin/sh*), and not the Bourne *sh*, *ksh*, or *csh*. In some cases, these files must also be read, and possibly modified by other scripts or the SAM program. For this reason, each variable definition must appear on a separate line, in the syntax:

```
variable=value
```

No trailing comments may appear on a variable definition line. Comment statements must be on separate lines, with the “#” comment character in column 1. An example of the required syntax for configuration files is given below.

```
# Cron configuration. See cron(1m)
#
# CRON: Set to 1 to start cron daemon
#
CRON=1
```

The name of a configuration script in */etc/rc.config.d* corresponds to the names of the associated startup and shutdown scripts found in */sbin/init.d*. In many cases the same 10-character base name may be used for the execution script, the sequencing links, and the configuration file. Here is an example using the cron subsystem.

<i>/sbin/init.d/cron</i>	execution script
<i>/etc/rc.config.d/cron</i>	configuration file
<i>/sbin/rc2.d/S730cron</i>	start sequence symbolic link
<i>/sbin/rc1.d/K270cron</i>	kill sequence symbolic link
<i>/usr/sbin/cron</i>	cron daemon (binary file)

The file */etc/TIMEZONE* contains the definition of the *TZ* environment variable. Its location is defined by standards documents. It is sourced by */etc/rc.config* along with the other */etc/rc.config.d/** files. The file */etc/rc.config.d/LANG* contains the National Language Support language specification.

A more detailed listing of */etc/rc.config.d/* configuration files is given in Appendix 2, “Configuration Files—HP-UX 10.20 Release.”

Sequencing Scripts with Link Files

The third aspect of the startup/shutdown design is controlling the order of script execution with link files. An important feature enables developers to control individual subsystems among run levels. This is accomplished by providing symbolic link files in run level directories that point to the scripts in */sbin/init.d*.

Run level Directories: */sbin/rc#.d*

The */sbin/rc#.d* (where # is a run level [0..6]) directories are startup and shutdown sequencer directories. They contain only symbolic links to startup/shutdown scripts in */sbin/init.d* that are executed by */sbin/rc* on transition to a specific run level. For example, the */sbin/rc3.d* directory contains symlinks to scripts that are executed when entering run level 3. (There is more information on */sbin/rc* below in “Run Levels and */sbin/rc*.”)

These directories contain two types of link files: start links and kill links. Start links have names beginning with the capital letter “S” and are invoked with the “start” argument at system boot time or on transition to a higher run level. Kill links have names beginning with the capital letter “K” and are invoked with the “stop” argument at system shutdown time, or when moving to a lower run level.

Further, all link files in a sequencer directory are numbered to ensure a particular execution sequence. Each script has as part of its name a three-digit sequence number. This, in combination with the start and kill notation, provides all the

information necessary to start up and shut down a system.

Figure 3 shows the run level directories and the relationship of the link files to the scripts. Because each script in */sbin/init.d* performs both the startup and shutdown functions, each will have two links pointing towards the script from */sbin/rc*.d*: one for the start action and one for the stop action.

Naming Conventions

In the naming conventions for the link files, the various components have the following meanings:

1. **Run Level Number:** The sequencer directory is numbered to reflect the run level for which its contents will be executed. In this case, Start scripts in this directory will be executed upon entering run level 2 from run level 1, and Kill scripts will be executed on entering run level 2 from run level 3.
2. **Sequencing Type:** The first character of a sequencer link name determines whether the script is executed as a start script (if the character is "S"), or as a kill script (if the character is "K").
3. **Sequence Number:** A three-digit number is used for sequencing scripts within the sequencer directory. Scripts are executed by type (start or kill) in lexicographical order.
4. **Script Name:** Following the sequence number is the name of the startup script. This name must be the same name as the script to which this sequencer entry is linked. In the example in Figure 3, the link points to */sbin/init.d/cron*.

Scripts are executed in lexicographical order. The entire file name of the link is used for ordering purposes. When adding new sequencer entries, sequencer numbers are chosen to allow for gaps so that future entries may be inserted without requiring renumbering of existing entries. (There is more information on sequence numbers below in "Link File Sequence Number Assignment.")

Subsystems are killed in the opposite order to their starting order. This implies that kill scripts will generally not have the same numbers as their start script counterparts. For example, if two subsystems must be started in a given order due to dependencies (e.g., *S111sys1* followed by *S222uses_sys1*), the counterparts to these scripts must be numbered so that the subsystems are stopped in the opposite order in which they were started (e.g., *K555uses_sys1* followed by *K777sys1*).

Also, kill scripts for start scripts in directory */sbin/rcN.d* reside in */sbin/rc(N-1).d*. For example, */sbin/rc3.d/S123system2*

and */sbin/rc2.d/K654system2* might be start/kill counterparts.

Core components of HP-UX will continue to support short filenames (i.e., 14 characters). Because of the filename placeholders for "K," "S," and the 3-digit sequence numbers, subsystem script names in */etc/init.d* should be restricted to 10 characters in length.

Run Levels and */sbin/rc*

In previous HP-UX releases, */etc/rc* was run only once. Now */sbin/rc* may be invoked several times during the execution of a system, once on each occasion the init process changes run levels. */sbin/rc* sequences through the execution scripts as it transitions through the run levels. However, only the subsystems configured for execution, through configuration variables in */etc/rc.config.d*, are started or stopped when transitioning the run levels.

Upon transition from a lower to a higher run level, the start scripts for the new run level and all intermediate levels between the old and new level are executed. Upon transition from a higher to a lower run level, the kill scripts for the new run level and all intermediate levels between the old and new level are executed.

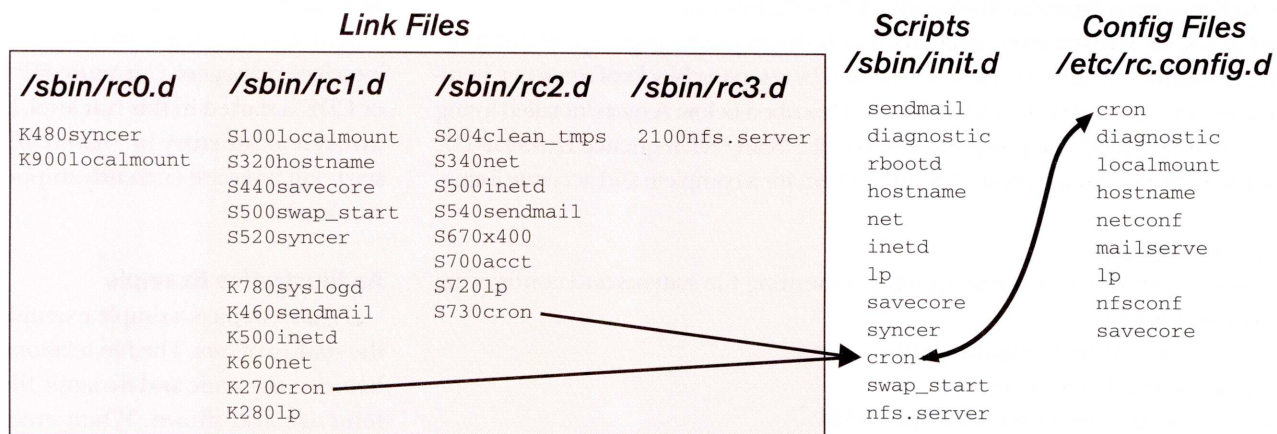
When a system is booted to a particular run level, it will execute startup scripts for all run levels up to and including the specified level (except run level 0). For example, if booting to run level 4, */sbin/rc* looks at the old run level (S) and the new run level (4) and executes all start scripts in states 1, 2, 3, and 4. Each level is sorted and executed separately to ensure that the lower level subsystems are started before the higher level subsystems.

Consequently, when shutting down a system, the reverse takes place. The kill scripts are executed in lexicographical order starting at the highest run level and working down, so as to stop the subsystems in the reverse order they were started. As mentioned earlier, the numbering is reversed from the startup order.

States 0 and S are special cases. When entering state 0, */sbin/rc* will run start scripts in */sbin/rc0.d*. Start scripts in state 0 are quick system administration scripts that prepare the system for a shutdown. When entering state 0 from a higher run level, the system is halted. When entering state S on startup, the init process executes entries only of type "sysinit" in */etc/inittab* and does not invoke */sbin/rc*. Instead, init execs a shell at the system console and the system is in single-user state. Table 1 summarizes the run level definitions.


Continued on Page 44

FIGURE 3 Startup/shutdown Component Relationships



Note: The sequence numbers above are for example only and may not accurately represent your system.

FIGURE 4 Implementing cron in the New Design



Link File Sequence Number Rationale and Assignment

The HP-UX 10.X operating system and its associated products group related functions into the same run state. This is accomplished by reserving blocks of sequence numbers for these functions. These assignments are described below. A more detailed listing of sequence numbers may be found in Appendix 1, "Start/Kill sequence Links HP-UX 10.20 Release." Please consult an HP-UX 10.X system for a complete and accurate listing.

Run Level 1 Paradigm

Run level 1 provides core services such as mounting file systems and configuring key system parameters.

0XX	reserved for temporary links
1XX	mount local file systems
2XX	essential process initialization/kill
3XX	set essential system parameters (hostname, LAN address)
4XX	set other system parameters (date, privilege groups)
5XX	start essential daemons (swapper and syncer daemons)
6XX-8XX	not currently used
9XX	reserved for future expansion

Run Level 2 Paradigm

Run level 2 is the general multi-user run state where most services are started.

0XX	reserved for temporary links
1XX	software installation/configuration (SD)
2XX	essential local daemons and services, started before network startup (clean log/tmp files, syslogd)
3XX	network startup
30X	network tracing/logging must be first
31X-33X	network link-level services (FDDI, ATM, Fiber, token ring)
34X	TCP/IP initialization (ifconfig, route, gateway, netmask, etc.)
35X-39X	other network startup (x25, loopback daemon, naming daemon)
4XX	NFS/NIS initialization
5XX-6XX	services built on top of network services (DCE, DFS, NCS, rbootd, NetLS, mail, etc.) (Also client/server services: X font server, Kanji server)
500	inetd super-server
7XX-8XX	other local daemons/services (lp, cron, diagnostics, auditing, accounting, etc.)
9XX	other links
900	"Don't Care" number for run state 2

Run Level 3 Paradigm

Run level 3 is the networked multi-user state. This run level is used to export file systems. Currently HP supports NFS exports.

0XX	reserved for temporary links
1XX	NFS exports (NFS server)
2XX-8XX	available
9XX	reserved for future expansion

Run Level 4 Paradigm

Run level 4 is reserved for graphical interface managers. Currently, HP VUE or CDE is started in this run level, but is invoked as an entry in */etc/inittab*. No start/kill links are currently shipped in run level 4.

An Illustrative Example

Figure 4 depicts a simple example for the startup of *cron*. The file relationships between the static and dynamic file systems are also shown. When entering run state 2 from a lower level, the "S" scripts are executed. In the example, *S730cron* is a link to the cron script under */sbin/init.d*. Cron will start because the configuration variable in */etc/rc.config.d/cron* is set to 1. A value of 0 would not start cron.

Current Assignments

Assignments for Start/Kill Sequence Links and Configuration Files for HP-UX Release 10.20 are given in the appendices to this article.

Startup Display on System Console

The 10.X startup display has adopted the screen-oriented interface used in the HP-UX 9.x Instant Ignition product. As individual subsystems are started, a status line is written to the system console; this is the string returned when the execution script is invoked with *start_msg* or *stop_msg*. The execution script returns a value of 0, 1, or 2 (or 3 for the special case of Reboot) which is displayed as OK, FAIL, or N/A. For terminals supporting HP terminal escape sequences, the table is displayed in screen-addressing mode; for other terminals the display is written in line (scroll) mode. A busy/wait status is

TABLE 1 *Run Level Definitions*

Run level	State	Sequencer Dir	/sbin/rc interaction
0	Halted	/sbin/rc0.d	All start and kill scripts executed
S	Single User		
1	Minimal System Configuration	/sbin/rc1.d	When entering from lower state, all start scripts are executed. When entering from higher state, all kill scripts are executed.
2	Multi-User	/sbin/rc2.d	
3	Exported File Systems	/sbin/rc3.d	
4	HP-VUE or CDE	/sbin/rc4.d	
5,6	Available	/sbin/rc5.d /sbin/rc6.d	

flashed if a subsystem takes longer than 5 seconds to start up. An example of a start-up display is as follows:

```
HP-UX Start-up in progress . . . . .Status

Mount file systems . . . . .[ OK ]
Setting hostname . . . . .[ OK ]
Set privilege group . . . . .[ OK ]
Display date . . . . .[ OK ]
Save system core image if needed . . . . .[N/A ]
Enable auxiliary swap space . . . . .[FAIL] *
Start syncer daemon . . . . .[ OK ]
```

* - An error has occurred !

* - Refer to the file /etc/rc.log for more information.

Startup Logging

All startup messages are logged to the file */etc/rc.log*. All execution script output should be directed to *stdout* or *stderr*, and not to */dev/console*. The log file */etc/rc.log* is rewritten on every system boot, so that it contains the history of all run-state transitions for the current bootstrap. The immediately previous log is preserved in the file */etc/rc.log.old*.

Extending the Startup Model

Adding Subsystems

System administrators and software product developers may need to add their subsystem(s) to this model. This may be accomplished easily by following a few steps.

A good place to start is with the execution script. This may be derived from any of the system execution scripts found in */sbin/init.d*. There is also a template in */sbin/init.d/template*. The guidelines for execution scripts must be strictly followed such that the interface with */sbin/rc* is correctly maintained. Any failure to do so will likely cause the execution script to fail.

A configuration file may also need to be added to */etc/rc.config.d*. This file should contain a variable assignment that enables the user or system administrator to enable or disable the particular subsystem. The file should also contain any other

variable assignments required by the subsystem during startup. The configuration file should not contain any executable script code other than variable assignments.

Selecting a run state and a sequence number should probably be done last. Many times it is difficult to determine exactly where the subsystem should fit until it has been coded and tested.

Selecting Sequence Numbers

All of the sequence numbers for HP products have been carefully assigned to ensure correct ordering, eliminate overlap, and allow room for growth. Each of these products has been registered with a central organization and has been given specific sequence numbers based on a number of factors about the product. When choosing sequence numbers for products or applications, system administrators and software product developers should adhere to the guidelines below.

Commercially available products may require specific sequence number ordering. Software developers should not arbitrarily choose an unused/unassigned number that happens to be absent from the system or table they are reviewing. Absent numbers that appear unused may have already been assigned by Hewlett-Packard to a product not yet

released or not installed on a particular system. Developers are encouraged to contact HP during their development cycle to determine correct assignments.

If specific ordering is not a concern and the subsystem may be started at any point after system boot and initialization, run level 2, number 900 should be used for the start link and run level 1, number 100 for the kill link. These "Generic" numbers may be used by any product or application without registering with Hewlett-Packard. Even if there are multiple entries of this type, the links are still unique because the entire filename is significant in ordering the scripts.

Guidelines for Startup/Shutdown Scripts

Startup execution scripts are installed in the `/sbin/init.d` directory and sequencer links are installed directly into the `/sbin/rcN.d/` directories directly from the install media. The startup execution scripts and the sequencing links are not customer editable. Administrators and users should remain aware of the following:

- Execution scripts must not be modified. If a script in `/sbin/init.d` is modified, the modification will be lost at the next update, when the script is overwritten with a new version.
- It is not necessary to remove scripts from the sequencer. If a script is removed from the sequencer, it will be replaced at the next update. Control variables in files within `/etc/rc.config.d` should be used to control whether a startup script is executed.
- Sequencer links must not be renamed. If a symbolic link in the sequencer directory is renamed (or renumbered), at the next update, a symbolic link with the original name will be installed. This will result in two copies of the same startup script appearing in the sequencer.

All HP-supplied startup scripts are meant to be present in sequencer directories as installed. They should not be removed or renamed (to change the relative sequencing).

Developing Startup/Shutdown Scripts

Developers should take care to design and test carefully the start and kill scripts for their products. A startup script that may work on one individual test system may cause intermittent or misleading failures on other systems that are configured differently unless the script has been robustly developed and tested. Here are several considerations a startup script developer should review when designing a script.

1. The script should be written for the POSIX shell (`/sbin/sh`), and not the Bourne shell, K-shell, or C-shell. The POSIX shell is the only shell that is guaranteed to be present on the root file system at system bootup.
2. Scripts should be carefully tested by the developer before they are installed on a system. Test by executing the script and explicitly reviewing the script input parameters and exit values. Scripts may be debugged by interpreting the script with the shell configured with the `-x` or `-v` option to display each line of the script as it is executed. This may be done easily by executing the command:

```
% /sbin/sh -v <your_startup_script> <input_parameters>
```

where `<your_startup_script>` is the name of your script and `<input_parameters>` may have the possible values (`start_msg`, `start`, `stop_msg`, `stop`). For each of these invocations of the startup script under test, the script exit values should be verified to be one of the allowed exit values: (0, 1, 2, 3). Please refer to earlier sections of this article for definitions of these input parameters and exit values.

3. Sequence numbers should be chosen carefully not to conflict with sequence numbers that may already be assigned. Even if a sequence number is not installed on your system in the directory `/sbin/rc[0-4].d/`, it may still be assigned to another product. Please refer to documentation in the directory `/usr/share/doc` for more complete listings of assigned startup sequence numbers.
4. Startup scripts must be written to use commands and libraries that are guaranteed to be present on the root file system of the system at bootup time. Startup scripts must not invoke any commands or libraries that are located under the `/usr`, `/opt`, or `/var` directories. In a multiple disk configuration, these directories are often located under file system mount points that are not available until late in the bootup cycle.
5. Startup scripts should be written to be portable across multiple product lines (S700 and S800) and as independent as possible of system configuration. A developer should test each startup script across a broad variety of systems and configurations.
6. Startup scripts should be written to work correctly when HP-UX is booted into any of the possible run-levels: (S, 1, 2, 3, 4). A developer should make sure that each script continues to work correctly when HP-UX is brought up

TABLE 2 *Startup File Naming and Packaging*

Filename	Purpose	Product	Fileset
/sbin/init.d/lp	Execution script	PrinterMgmt	LP-SPOOL
/sbin/rc1.d/K280lp -> /sbin/init.d/lp	Start sequence link	PrinterMgmt	LP-SPOOL
/sbin/rc2.d/S720lp -> /sbin/init.d/lp	Kill sequence link	PrinterMgmt	LP-SPOOL
/etc/rc.config.d/lp	Configuration file	PrinterMgmt	LP-SPOOL
/usr/sbin/lpsched	Daemon (binary)	PrinterMgmt	LP-SPOOL
/usr/bin/lp	Command (binary)	PrinterMgmt	LP-SPOOL

in each of these run-levels and transitioned between various combinations of these run-levels.

File Name Conventions, Permissions, and Fileset Packaging

File name conventions have been presented in previous sections. Often, a developer may use the same 10-character file basename for the various startup file names and for associated daemons. Execution scripts are typically delivered with (owner, group, permissions) = (bin, bin, 0555 read/execute), configuration files with permissions = 0444 read. All startup and configuration files should be delivered in the same fileset as the rest of the product, so that startup files can be installed and uninstalled at the same time as the product is installed. This approach helps prevent invalid and potentially incorrect startup files from remaining on the system after a product is removed or updated. An example of file naming and fileset packaging is shown in *Table 2*.

Developer Assistance

All of the sequence numbers for HP products have been carefully assigned to ensure correct ordering, eliminate overlap, and allow room for growth. Each of these products has been registered with a central organization and has been given specific sequence numbers based on a number of factors about the product. When choosing sequence numbers for products or applications, system administrators and software product developers should adhere to these guidelines.

Commercially available products may require specific sequence number ordering. Software developers should not arbitrarily choose an unused/unassigned number that happens to be absent from the system or table they are reviewing. Absent numbers that appear unused may already have

been assigned by Hewlett-Packard to a product that has not yet been released or is not installed on the system. Instead, developers should contact their Hewlett-Packard representative.

Developers will be asked to complete a checklist about their product and assigned a set of numbers (start and kill). The product will then be registered with Hewlett-Packard and be ensured that no other registered products will have duplicate numbers that may corrupt an end user's system. However, duplicate numbers may be assigned for applications that do not conflict. Developers are encouraged to contact HP late in their development cycle, as close to application release as possible.

Summary

This article has reviewed the changes to the HP-UX startup and configuration models. Users are encouraged to review the documentation delivered with the HP-UX 10.X releases under the directory */usr/share/doc* for the most recent changes and assignments to this system. Developers seeking additional assistance should contact their Hewlett-Packard representative.

The following individuals contributed to this project and article: Edgar Circenis, Frank Feather, Dave Gutierrez, Jeff Krenek, Bruce Rodean, Todd Poynor, and Ron Smith. ■

John Fenwick is a development engineer in Hewlett-Packard's Open Systems Software Division in Cupertino, California.

Appendix 1. Start/Kill Sequence Links-HP-UX 10.20 Release

Start	Kill	Notes
rc1.d/S099mel200escs		extended SCSI driver
rc1.d/S100localmount	rc0.d/K900localmount	mount local filesystems
rc1.d/S140upgrade		9.0->10.0 migration tools
	rc0.d/K800killall	stop processes on shutdown
rc1.d/S300switchover	rc0.d/K700switchover	Switchover high availability
rc1.d/S320hostname		system name
rc1.d/S400set_privgrp		privilege groups
rc1.d/S420set_date		time/date
rc1.d/S440savecore		save system core images
rc1.d/S500swap_start		configure swap space
rc1.d/S520syncer	rc0.d/K480syncer	syncer daemon
rc2.d/S008net.sd		SD/DCE
rc2.d/S100swagentd	rc1.d/K900swagentd	SD - install
rc2.d/S120swconfig		SD - configure
rc2.d/S200clean_ex		200: clean files
rc2.d/S202clean_uucp		
rc2.d/S204clean_tmpps		
rc2.d/S206clean_adm		
rc2.d/S220syslogd	rc1.d/K780syslogd	system logging daemon
rc2.d/S230ptydaemon	rc1.d/K770ptydaemon	pty allocator daemon
rc2.d/S300nettl	rc1.d/K700nettl	300: netwk start: trace/log daemon
rc2.d/S310atm		netlink ATM
rc2.d/S312fddi		netlink FDDI
rc2.d/S314fiber		netlink Fiber Channel
rc2.d/S316token		netlink Token Ring
rc2.d/S318vgal		netlink 100VG Anylan
rc2.d/S320hpether		Ethernet drivers
rc2.d/S320hpgv		100VG driver for NIO S800s
rc2.d/S322hpbt		100BT driver for NIO S800s
rc2.d/S330hippi	rc1.d/K670hippi	HIPPI networking
rc2.d/S340net	rc1.d/K660net	340: net start
rc2.d/S346acc	rc1.d/K654acc	Native ACC (X25)
rc2.d/S350x25		X25 networking, ISU product
rc2.d/S352isdn	rc1.d/K648isdn	ISDN Transport
rc2.d/S370named	rc1.d/K630named	Internet domain name server
rc2.d/S380ots	rc1.d/K620ots	OSI Transport Services
rc2.d/S390nwspport	rc1.d/K610ots	Netware Transport Services
rc2.d/S400nfs.core	rc1.d/K600nfs.core	400: nfs/nis/nis-plus services
rc2.d/S406nisplus.server	rc1.d/K594nisplus.server	
rc2.d/S408nisplus.client	rc1.d/K592nisplus.client	
rc2.d/S410nis.server	rc1.d/K590nis.server	
rc2.d/S420nis.client	rc1.d/K580nis.client	
rc2.d/S430nfs.client	rc1.d/K570nfs.client	
rc2.d/S490mrouted	rc1.d/K510mrouted	Multi-casting routing daemon
rc2.d/S500inetd	rc1.d/K500inetd	500: netwk svcs: inet super-server
rc2.d/S510gated	rc1.d/K490gated	gateway routing daemon
rc2.d/S520rdpd	rc1.d/K480rdpd	router discovery protocol daemon
rc2.d/S522ppp	rc1.d/K478ppp	Point to Point Protocol
rc2.d/S525rarpd	rc1.d/K475rarpd	RARP protocol
rc2.d/S530rwhod	rc1.d/K470rwhod	remote-who daemon
rc2.d/S540sendmail	rc1.d/K460sendmail	Internet mail delivery
rc2.d/S545news	rc1.d/K475news	news services
rc2.d/S550ddfa	rc1.d/K450ddfa	DTC terminal servers
rc2.d/S560SnmpMaster	rc1.d/K440SnmpMaster	SNMP network management
rc2.d/S565SnmpHpunix	rc1.d/K435SnmpHpunix	network management sub-agents
rc2.d/S565SnmpMib2	rc1.d/K435SnmpMib2	
rc2.d/S5650OspMib	rc1.d/K435OspMib	
rc2.d/S566cmsnmpagt	rc1.d/K434cmsnmpagt	Service Guard SNMP daemon.
rc2.d/S570dce	rc1.d/K430dce	Distributed Computing services
rc2.d/S580dfs	rc1.d/K420dfs	Distributed Filesystem
rc2.d/S590ncs	rc1.d/K410ncs	NCS broker daemon

Appendix 1. Start/Kill Sequence Links-HP-UX 10.20 Release, *continued*

Start	Kill	Notes
rc2.d/S600i4lmd	rc1.d/K400i4lmd	NetLS network license server
rc2.d/S610rbootd	rc1.d/K390rbootd	remote boot protocol
rc2.d/S620xfs	rc1.d/K380xfs	X font server
rc2.d/S630vt	rc1.d/K370vt	virtual terminal server daemon
rc2.d/S640kks	rc1.d/K360kks	Kanji server
rc2.d/S642egcd	rc1.d/K358egcd	Kanji server
rc2.d/S644vjed	rc1.d/K356vjed	Kanji server
rc2.d/S646jsrver	rc1.d/K354jsrver	Kanji server
rc2.d/S650dtcmgr	rc1.d/K350dtcmgr	DTC manager
rc2.d/S660xntpd	rc1.d/K340xntpd	network time synchronization
rc2.d/S670x400	rc1.d/K330x400	X400 OSI Messaging
rc2.d/S680snaplus		SNAPlus
rc2.d/S690nwsrvr	rc1.d/K310nwsrvr	Netware File/Print Services
rc2.d/S700acct	rc1.d/K300acct	system accounting
rc2.d/S710hparrray	rc1.d/K290hparrray	disk arrays
rc2.d/S720lp	rc1.d/K280lp	UNIX print services
rc2.d/S722pd	rc1.d/K278pd	Distributed Print services
rc2.d/S730cron	rc1.d/K270cron	timed-job execution daemon
rc2.d/S740supprtinfo		system diagnostics
rc2.d/S742diagnostic	rc1.d/K258diagnostic	system diagnostics
rc2.d/S750envd	rc1.d/K250envd	environmental monitoring daemon
rc2.d/S760auditing	rc1.d/K240auditing	system auditing
rc2.d/S770audio	rc1.d/K230audio	audio server
rc2.d/S780cmcluster	rc1.d/K220cmcluster	disk mirroring
rc2.d/S790rti		real-time interface
rc2.d/S800spa		NFSD/SPA admin.
rc2.d/S805orbplus	rc1.d/K195orbplus	ORB Plus obj_locator
rc2.d/S810scope	rc1.d/K190scope	Performance tools: scopeux
rc2.d/S810mwa	rc1.d/K190mwa	Perf.: measurement works agent
rc2.d/S814pv	rc1.d/K186pv	Perf.: PerfView monitor
rc2.d/S820prm		Process Resource Manager (FSS)
rc2.d/S824rbt	rc1.d/K176rbt	Convex Robotics server
rc2.d/S830opcagt	rc1.d/K170opcagt	Operations Center agent
rc2.d/S834asu	rc1.d/K166asu	Advanced Server for UNIX (NT)
rc2.d/S835opckmsg	rc1.d/K165opckmsg	Operations Center message log
rc2.d/S838omni	rc1.d/K162omni	Omniback system backup
rc2.d/S840dkit	rc1.d/K160dkit	Datakit connectivity
rc2.d/S850ns-commerce	rc1.d/K150ns-commerce	Netscape Commercial Server
rc2.d/S851ns-communic	rc1.d/K149ns-communic	Netscape Communications Server
rc2.d/S852ns-commerce	rc1.d/K148ns-commerce	Netscape Proxy Server
rc2.d/S854omi	rc1.d/K146omi	Open Market Server
rc2.d/S855raptor	rc1.d/K145raptor	Raptor Firewall
rc2.d/S880swcluster		Raptor Firewall
rc2.d/S900<filename>	rc1.d/K100<filename>	Don't care slot for Runstate 2
rc2.d/S900psafe	rc1.d/K100psafe	Ebacom Powersafe 5 env. monitor
rc2.d/S940OV500	rc1.d/K060OV500	OpenView product
rc2.d/S980upgrade		9.0>10.0 migration tools
rc3.d/S100nfs.server	rc2.d/K900nfs.server	100: NFS export services
rc3.d/S300tbroker	rc2.d/K700tbroker	TaskBroker
rc3.d/S800adc.server	rc2.d/K200adc.server	AdminCenter
rc3.d/S810chemserver	rc2.d/K190chemserver	Chemserver data collection/analysis
rc3.d/S900<filename>	rc2.d/K100<filename>	"Don't care" slot for Runstate 3
rc3.d/S990dtlogin.rc	rc2.d/K100dtlogin.rc	CDE environment

Appendix 2: Configuration Files–HP-UX 10.20 Release

rc.config file	init.d script	Fileset	Notes
LANG	–	UX-CORE	default language setting, use C if unset
TIMEZONE	/sbin/rc	UX-CORE	location is /etc/TIMEZONE
set_date	set_date	UX-CORE	date format
cron	cron	UX-CORE	cron daemon configuration
envd	envd	UX-CORE	environmental monitoring daemon
savecore	savecore	UX-CORE	savecore save directory
syncer	syncer	UX-CORE	syncer on/off
list_mode		UX-CORE	/sbin/rc checklist list modes
clean	clean_adm clean_ex	UX-CORE	list/remove file options
clean_tmps	clean_tmps	UX-CORE	list/remove /tmp files
clean_uucp	clean_uucp	UUCP	list/remove uucp tmp files
lp	lp	LP_SPOOL	lp spooler configuration
dp	dp	DistributedPrint	Distributed Print spooler configuration
acct	acct	ACCOUNTNG	system accounting
audio	audio	AUDIO-SRV	audio subsystem
auditing	auditing	AUDIT	system auditing
netconf	hostname net rdpd rarpd	InternetSrvcs	hostname, IP address, netmask
nettl	nettl	NETTL-RUN	network tracing/logging
hpetherconf	hpether	Networking	ethernet link drivers
hpfddiconf	fddi	Networking	FDDI link drivers
hptokenconf	token	Networking	Token-ring link drivers
i4lmd	i4lmd	LSSERV-RUN	Formerly netls, now i4lmd - "Gradient"
nfsconf	nfs.client nfs.server	NFS_RUN	NFS, automount configuration
namesvrs	named nfs.client nfs.server nis.client nis.server	INETSVCs-RUN NFS-CLIENT NFS-SERVER NIS-CLIENT NIS-SERVER	NFS services NIS services
netdaemons	rwhod xntpd rbootd	InternetSrvcs	rwho daemon time syqchronization remote boot daemon rbootd
mailservs	sendmail	InternetSrvcs	mail configuration/addressing
SnmpMaster	SnmpMaster	MASTER	SNMP product - master daemon
SnmpHpunix	SnmpHpunix	SUBAGT-HPUNIX	SNMP product - HP-UX sub-agent
SnmpMib2	SnmpMib2	SUBAGT-MIB2	SNMP product - sub-agent
cmsnmp	cmsnmp	ServiceGuard	Service Guard SNMP reporting service
swconfig	swconfig	SW-DIST	SD Distribution Tools
vt	vt	SYSCOM	HP Virtual Terminal
ptydaemon	ptydaemon	CMDs-AUX	pty allocator configuration
x25	x25 dtc	X25	X.25 configuration DTC Terminal controller
dce	dce	DCE-CORE	DCE services
dfs	dfs	DFS-CORE	DFS file system
ncs	ncs	NCSNCK	NCS Location Broker
ots	ots	OTS-RUN	OSI Transport Services–ISU
scope	scope	SCOPE	Performance Tools: scopeux–ISU

ISAMATION

THE FASTEST ISAM REBUILD TOOL AVAILABLE!

**Go fishing, play golf, or watch TV
but don't spend your weekend in the
office reorganizing data files!**



*For more information on the ultimate ISAM
maintenance and data recovery tools, contact:*

**Computer Solutions, Inc.
of Orlando**

120 E. Marks St. • Suite 225 • Orlando, FL 32803


407-649-0123 or 512-343-6634

FAX 407-649-1407 email - info@csisoft.com

CIRCLE 154 ON READER SERVICE CARD

Appendix 2: Configuration Files-HP-UX 10.20 Release, *continued*

rc.config file	init.d script	Fileset	Notes
diagnostic	diagnostic	SUP-CORE-700	Diagnostics-ISU
supprtinfo	diagnostic	SUPPORT-INFO	Diagnostics-CSU
rdpd	rdpd	NET-RUN	Router Discovery protocol
spa	spa	SystemAdmin	AM/NFS Diskless configuration/startup
rti	rti	DART	Real time interface
x400	x400	DART	X400 OSI Messaging
cmcluster	cmcluster	CM-CORE	Disk mirroring
switchover	switchover	SWITCHOVER-RUN	Starfish project
hparray	hparray	C2400-UTIL	Disk Arrays - checks disk array parity
xfs	xfs	X11-FONTSRV	X11 font server
hpetherconf	hpetherxfs	LAN-RUN	Ethernet LAN configuration (ethernet address)
prm	prm	PRM-RUN	Process Resource Manager (Fair Share Sched.)
hpvg	hpvg	Networking	100VG networking
hpbt	hpbt	Networking	100BT networking
hpgalconf	hpgal	Networking	100VG networking
asu	asu	ASU.ASU	Advanced Server for UNIX (NT services)
rbt	rbt	RBTD-BASE	Convex Robotics server
dkit	dkit	DKIT-RUN	DataKit Connectivity product
hippi	hippi	HIPPI-RUN	HIPPI network services
jserver	jserver		Kanji server
kks	kks	IMX11-JPN-RUN	Kanji server
egcd	egcd	EGB-RUN	Kanji server
vjed	vjed	VJE-RUN	Kanji server
ovlmd	ovlmd		OpenView SNMP product



SLIP on HP-UX

SLIP (Serial Line Internet Protocol) is the poor relation of networking. Instead of 10 megabit per second Ethernet cabling or 100 megabit per second fiber optic cabling, the superhighways, or even the 56K or 128K of frame relay or dedicated lines, you are using plain old serial cable and ordinary modems. You are on the secondary roads the county doesn't pave too often. (I promise that is the last information superhighway metaphor.) You can use SLIP with POTS (Plain Old Telephone Service) with current state of the art 28.4 baud or as slow as whatever modem you have. No one uses SLIP if a faster connection is available. Because of this, not everything for SLIP is set up for you when you unwrap your HP box for initial install. Setting up SLIP is a little bit of a black art and is a continual source of questions on the HP newsgroups.

Why SLIP?

Since SLIP is slower than Ethernet LAN and requires more effort to set up, why bother? No one would bother if they had an alternative. If every computer you want to network to is connected to you by a chain of high-speed cable, you can lead a full and happy life without SLIP. Of course, you may be using SLIP without knowing it. Your connection to another machine may be Ethernet, but the second machine's to a third may involve SLIP somewhere down the line. You don't care—it looks like every other network connection; only the system administrators at the individual sites are concerned with the hardware. Well, you may care if the speed of the connection becomes noticeable. When your upload or display is moving at glacial slowness, you can justly suspect that there is a SLIP connection somewhere in the chain.

The "don't know, don't care" is the powerful argument for SLIP. The network-aware applications, rsh, telnet, ftp, sockets, NFS, X, and all the browsers don't care about the mechanics of the network connection. They sit layers above this and will function happily over a SLIP line. They may be happy, but if they use a lot of bandwidth, you the user may not be happy with the response time. In particular, X using a lot of graphics and NFS will not result in happy campers. This limitation using a skinny wire with all the bits in a row is the driving force behind special adaptations such as Low Bandwidth X.

by Larry Headlund

Illustration by Stephen Kramer

If you're an HP Computing Professional, you may be missing an important connection.

If you aren't a member of Interex, the International Association of Hewlett-Packard Computing Professionals, you're missing out on a number of valuable connections.

- Publications written to help you solve HP-UX problems
- Conferences designed to increase your knowledge and competitiveness
- Special Interest Groups to connect you with other users
- Advocacy efforts that give you a way to protect your investment in HP equipment

Get connected!

Find out more about these valuable resources by calling us at **800.INTEREX** or **408.747.0227**

Interex Online
<http://www.interex.org/>

interex

*The International
Association of
Hewlett-Packard
Computing Professionals*

Now that I have mentioned browsers, I touch on what is the agenda for most of the interest in SLIP, particularly in the non-UNIX community. "Surfing the Web" is the carrot that motivates the donkey of all those SLIP programs on desktops. It is probably of more than passing interest to the readers of this article. All that multiple simultaneous access to and from your machine and all those socket programs famously sending e-mail to Afghanistan are all right, but they can't compete in glamour with pointing your Mosaic browser at somewhere interesting.

Some people use SLIP to host their own Web pages. This is fine as long as your page does not make heavy use of graphics and you don't expect a large volume of hits. If either condition does not hold, people accessing your page may experience unacceptably slow performance. You have to balance this relatively slow response against the costs of maintaining a high-speed link.

After listing the good and not so good reasons for SLIP, I must mention a bad one: transferring binary files. Some people feel they should implement SLIP so they can transfer non-ASCII files. SLIP is a very inefficient tool for this purpose. The reason is that the Internet Protocol has a lot of nice features, but efficiency is not one of them. The Internet Protocol was designed for fast connections and breaks up its information into packets with identifying information included. This allows the packets to be sent individually and reassembled at the destination. This allows multiple paths to a destination and retransmission until the communication is complete. On the downside it adds the overhead of that identification to every packet. This is relatively irrelevant when you are measuring bandwidth in megabits per second but important when you are dealing in tens thousands of Baud at the best. Protocols designed for serial communications such as those used in kermi or the (x)(y)(z) modem family are much more efficient. In fact, when you have a clean line and an ASCII file, the *cu* procedures are the most efficient for transferring files. With large files and a clean line, it can even be more efficient to compress the file, uuencode it, transfer the file with *cu*, and decode and uncompress it at the other end. Note I say with a clean line. With a noisy line the binary transfer capabilities of communication programs pay for themselves. Only when a consistent interface or ease of use can dominate transfer time is SLIP (ftp, rcp) the right choice.

SLIP and ppl

HP-UX implements SLIP with a program called *ppl* (for Point to Point Link). You get *ppl* with HP's LAN package,

therefore if you can do any networking you have it available. The *ppl* software is part of the NETINET file set. With *ppl* and supporting utilities you can have dial in access to your machine, dial out from your machine, and direct connection over serial lines. This last is most useful when you need temporary connection to a PC with its own SLIP program but no networking hardware.

In writing all this I referred frequently to the HP publication *Using Serial Line IP Protocols*. If you opted for the paper documentation with your system, it is included, a thin paper-bound volume. It comes with a three-hole punch so someone may have secreted it in with another networking volume. It would also be on your CD-ROM of all HP-UX manuals.

Before You Start

In all that follows I am going to refer to your local machine as A and the remote machine as B. Assuming you already have an Internet Address (nnn.nn.nn.nn) for A, you must have a different Internet address for A to use for SLIP. Let me repeat that. You must use a different address for SLIP. Using the same address, which would be the natural thing to do, will result in an interesting network crash. Interesting to an outside observer—those on the scene may use other adjectives. This address should also be on a subnetwork different from your normal one. You want this new subnetwork to be for the exclusive use of SLIP.

Suppose you have a network that uses 192.10.10 with no subnetting. Let us say your normal IP address for A is 192.10.10.1 and that B's IP address will be 192.10.10.2. The normal rules of avoiding address collision apply. For the SLIP address on A we will use 192.1.1.1. Note this is on a different subnet from 192.10.10. If you make the SLIP address on A 192.10.10.3, you will be sorry.

What is required next is to set up A, your machine, with this information. If B is not an HP-UX machine, you will use the procedures on B's SLIP package to have it use B's IP address 192.10.10.2 and to talk to A on 192.1.1.1.

If Both machines are HP-UX, then what follows will work on B also. The thing to keep straight is the addresses:

For A: 192.1.1.1 is local 192.10.10.2 is remote

For B: 192.10.10.2 is local 192.1.1.1 is remote.

etc/netlinkrc

This is the script that handles all the network startup (for HP-UX 9.x). You are modifying this file so that A, your host, will respond to ARPs (Address Resolution Protocol) for B.

Today's the day
your overworked
processor spontaneously
combusts.



Are you prepared?

CLAM Associates can help you prevent the consequences of downtime. Power failures and surges cause 45.3% of data loss, but making your mission-critical applications continuously available is simple with CLAM Integration Services. Buying the work for you is *our* part. We solutions; our expertise in cluster reliable source to integrate into your system. Call (617) 621- or surf <http://www.clam.com> for quality services will provide



anything can happen.

Associates' MC/ServiceGuard software is the easy part; making it specialize in continuous availability architecture makes CLAM the most Hewlett-Packard's MC/ServiceGuard 2542, e-mail marketing@clam.com more information on CLAM. Our peace of mind in a world where

CLAM Associates. Solutions for a strange, strange world.

CIRCLE 6 ON READER SERVICE CARD

First you need to get the LAN address by running *lanscan*. under "Station Address" you will see a number like 0x08000987E2B8. You add the following line to */etc/netlinkrc*:

```
arp -s 192.10.10.2 8000987E2B8 pub
```

This command will be run every time you boot up. You can run it from the command line during setup and testing.

Device Files

If you have not set up your computer for UUCP communications, you must do that first. You must be able at least to use *cu* to call out on the ports you are interested in. The setup for dialing in, dialing out, and hard-wired communication with SLIP uses the UUCP files *Systems* and *Devices*. Being able to *cu* out to a port is not only essential in debugging SLIP when you are setting it up, but is also required for some forms of SLIP itself. In particular, when you are connecting to a system that dynamically assigns IP addresses, giving you a different IP address every time you call out, you start the process by using *cu* to connect, start up the SLIP process on the other end, and use the IP address it gives you to modify manually some of your own files. I will go through this process in detail later; the point is you have to have *cu* working even to begin.

If you are starting from scratch, you have to make the device file for the port you are going to use. You could use the *mknod* command and edit the files */usr/lib/uucp/Devices* and */usr/lib/uucp/Systems* yourself. If you are uncertain about how to do this, I would recommend letting SAM do it for you. You can always edit the */usr/lib/uucp/** files yourself later if you need to. You also need a magic file that *ppl* uses to pass packets back and forth. This file is named */dev/ni* and is probably already on your system. If it isn't, issue the following command:

```
$ mknod /dev/ni c 56 0
```

ppl Configuration

ppl depends on three configuration files in */usr/lib/ppl*:

```
ppl.users
ppl.ipool
ppl.remotes
```

Sample files should have been put in this directory by the

installation process. If they are not there, I will give you sufficient samples later of each. The permissions and ownership of all three are the same:

```
-r--r--r-- 1 bin bin
```

ppl.users

This is a convenience file for dialing in or out. The program *ppl* has an argument identifying the remote host, either where you are dialing out to for dial out from your machine or where you are dialing in from for dial-in access. An entry in the *ppl.users* file means a given user has a default remote host and need not type it in. The format looks like this:

```
<username> <inetaddress or rhostname>
```

```
lmh 192.10.10.2
```

where *lmh* is the user's name and shown is the name of an entry in the */etc/hosts* file.

ppl.ipool

This is a file you only use if you have a modem pool set up for dial in. Essentially, *ppl.ipool* is just a list of Internet addresses (nnn.nn.nn.nn) that *ppl* can use and is used only if a specific port and local Internet are allocated for use by that host.

ppl.remotes

This is the file everyone needs for every connection. It defines the type of connection, the address, the serial line characteristics, everything. It has a flavor similar to the *uucp Systems* file and is similarly touchy. I will give examples for dial in with login, dedicated dial in, dial out with login, dial out without login, direct connection, and dialing out to an unknown remote address. This last is a very common situation, particularly when you are dealing with an Internet Service Provider. Life being what it is, it is also a situation which *ppl* does not handle elegantly.

In all the examples, the HP at your site, the local machine, will be called A and the remote machine will be called B. On A, A's local address is 192.1.1.1 and B's address is 192.10.10.2. *Listing 1* shows the format of *ppl.remotes*. In all the examples I will use SLIP as the protocol. PPP, the newer serial protocol, is listed as an option but not supported and ASLIPC (Abbreviated Serial Line Internet Protocol Client) and ASLIPS (Abbreviated Serial Line Internet Protocol Server) are not as popular.

LISTING 1 *ppl.remotes format*

```

# remote hostname or Internet address
# local hostname or Internet address
# Internet mask
# protocol [SLIP] [ASLIPC] [ASLIPS] [PPP]
# type [DIRECT] [DIALIN] [DIALOUT] [DIALIN & DIALOUT]
# UUCP system name
# line parity [EVEN] [ODD] [NONE]
# line speed
# serial line
# phone number
# modem control available
# log in info
# command name

```

Dial In with Login

In this scenario the outside user connects to your machine as normal. That is, there is a *getty* or *uugetty* running on the port, the user has an account and password, etc. The user would be running a communication program on his computer, B.

Sometime after logging in, the user just issues the shell command

```
$ ppl
```

If there is no entry in *ppl.users* for this user, he will have to give the hostname or address with the *ppl* command:

```
$ ppl 192.10.10.2
```

After the *ppl* command starts up on A, the user exits the communication program on B or whatever the procedures are on B. Many of the PC programs have utilities to do this automatically.

The appropriate *ppl.remotes* format on B is shown in *Listing 2*. Since the *getty/uugetty* program takes care of most of the communications setup for the line, you only need tell *ppl* what the parity is and if modem control is available.

The advantage of a setup like this is that the same incoming line can be used for shell access and for SLIP connections.

Dial in without Login

In this case there is no *getty* or *uugetty* running on the line. Only SLIP connections occur on the line. The *ppl* program takes care of all line control. Therefore, you have to add another line to the */etc/netlinkrc* file to start *ppl* at system boot. The line would look like

```
ppl -o 192.10.10.2 &
```

LISTING 2 *Dial in with login*

```

192.10.10.2 # remote hostname or Internet address
192.1.1.1   # local hostname or Internet address
            # Internet mask
SLIP        # protocol
DIALIN      # type
            # UUCP system name
NONE        # line parity [EVEN] [ODD] [NONE]
            # line speed
            # serial line
            # phone number
NO          # modem control available
            # log in info
            # command name

```

LISTING 3 *Dial in without login*

```

192.10.10.2 # remote hostname or Internet address
192.1.1.1   # local hostname or Internet address
            # Internet mask
SLIP        # protocol
DIRECT      # type
            # UUCP system name
NONE        # line parity [EVEN] [ODD] [NONE]
19200       # line speed
/dev/ttyd01 # serial line
            # phone number
NO          # modem control available
            # log in info

```


LISTING 4 *Dial out without login*

```

192.10.10.2 # remote hostname or Internet address
192.1.1.1   # local hostname or Internet address
            # Internet mask
SLIP        # protocol
DIALOUT     # type
bprime      # UUCP system name
NONE        # line parity [EVEN] [ODD] [NONE]
19200       # line speed
            # serial line
            # phone number
NO          # modem control available
            # log in info
            # command name

```

LISTING 5 *Dial out with login*

```

192.10.10.2 # remote hostname or Internet address
192.1.1.1   # local hostname or Internet address
            # Internet mask
SLIP        # protocol
DIALOUT     # type
            # UUCP system name
NONE        # line parity [EVEN] [ODD] [NONE]
19200       # line speed
            # serial line
19005551212 # phone number
NO          # modem control available
login: guest ssword: guest guest_start # log in info
"" ppl \n   # command name

```

LISTING 6 *Direct connection*

```

192.10.10.2 # remote hostname or Internet address
192.1.1.1   # local hostname or Internet address
            # Internet mask
SLIP        # protocol
DIRECT     # type
            # UUCP system name
NONE        # line parity [EVEN] [ODD] [NONE]
19200       # line speed
/dev/ttyd01 # serial line
            # phone number
NO          # modem control available
            # log in info

```

Listing 3 shows what the *ppl.remotes* file would look like for this. Note that line speed, parity, and the specific device used must be given.

If you are doing this, put a dummy line in the */etc/inittab* file to document that this line is dedicated to *ppl*. Perhaps:

```
c0:4:off:# ppl dial in on /dev/ttyd01
```

Dial Out without Login

Now we are getting to the good stuff, letting you cruise the Internet with the rest of the semi-civilized world. The following assumes the remote machine B has a dedicated SLIP line running and you always get the same IP address from B and that B knows A's IP address. In this form *ppl* uses the facilities of *uucp*. That is, you must have a valid entry in */usr/lib/uucp/Systems*, here with the name *bprime*, for connecting to B using *uucp*. The password for B, if any, and other information is in the */usr/lib/uucp/Systems* file. (See *Listing 4*.)

With this setup the user on A would type

```
$ ppl -o
```

If all goes well after the connection is established the user sees the message:

```
initialization complete.  protocol running
```

The user can then telnet, ftp, run the browser, whatever.

Dial Out with Login

Listing 5 illustrates this setup. The setup is much like that for dialing in without login but with the difference that the information on logging in, which would normally be contained in the */usr/lib/uucp/Systems* file, is instead in *ppl.remotes*. If you are only and always contacting remote system B for the purpose of running a SLIP connection, either way would work. If you sometimes reach out to B for shell access or *uucp* file transfers, then this is the preferred setup. The user issues the same commands as for dial out without login to make the connection.

Direct Connection

This is the case where the two boxes are at the same location and you have them hard-wired together without any modems.

The setup is simple since there are no modems to fool with. After adding something like Listing 6 to *ppl.remotes*, you would add the line

```
ppl -o 192.10.10.2 &
```

to your */etc/netlinkrc* script. Note the similarities to dial in without login.

Dialing Out to an Unknown Remote Address

Many Internet Service providers dynamically assign you a different IP address when you dial in. HP *ppl* does not have the facilities to read this address itself and do the appropriate configuration. You have to do some things manually. Basically, you will have to use *cu* or another communications program to get things rolling. This is yet another instance where it helps to have your *uucp* setup ready before you begin. You should also have a *ppl.remotes* entry for a direct connection (Listing 6) ready and waiting, but do not have the *ppl* command running from */etc/netlinkrc*.

Step one is dialing into the remote machine B. I will assume you are using *cu*. Probably you have to give a password and issue whatever command they give you to start up SLIP. This command on B will echo back an IP address.

Step two is to open another window on A and edit *ppl.remotes* and include this IP address. The security problems this opens up on your machine are left as an exercise for the reader. Step three is to start *ppl* in the *cu* window with the command

```
~&ppl -o <ipaddress>; sleep 9999
```

Don't Let Poor System Performance Drive You Crazy!



System managers today are being faced with complex problems, performance management is one of them. **SOS/9000 Performance Advisor** is the ultimate on-line performance tool for HP 9000 systems. With instant access to performance data as it happens, system managers can optimize performance, identify bottlenecks, increase throughput and solve performance problems quickly and easily.

For more information or a **FREE 30 day trial copy**, call today at: (541) 926-3800!



SOS/9000
Performance Advisor

LUNR
PERFORMANCE SOLUTIONS
Performance Beyond Expectation
240 2nd Ave. SW, Albany, OR 97321, USA

CIRCLE 129 ON READER SERVICE CARD

The *~&* takes over the input and output from *cu*.

You can now run any network commands from any window other than the one you dialed out on. Iconize that one until it is time to break the connection.

Conclusion

HP does have a SLIP implementation that works rather well. It does not handle the newer PPP protocol and it is a bit more difficult to administer than some products, and it does not adapt to dynamic IP addresses elegantly. However, it is right there if you want to use it.

I based all of the examples on HP-UX 9.x. At 10.x the *uucp* files *Systems* and *Devices* are in */etc/uucp* and the

three *ppl* files *ppl.remotes*, *ppl.users*, and *ppl.ipool* are in */etc/ppl*. The *netlinkrc* script has been split up into several files, each of which controls the startup and shutdown of a particular service. All are found in */sbin/init.d*. You can create a file for *ppl* using one of the existing scripts as a model. If you add code to an existing script, it may be overwritten at a subsequent upgrade. ■

Larry Headlund is the president of Mathematical Engineering, Inc., a UNIX and Motif development company, and has been working with commercial UNIX since 1983 and with X since 1987. He writes the X-Watch column for hp-ux/usr. He can be reached at lmh@world.std.com or at 1.617.242.7741.



by Judy Tarutz

IN THIS ISSUE WE PROFILE the Haystack Observatory, a radio-astronomy research facility of the Massachusetts Institute of Technology. We spoke with Dr. Alan Whitney, Associate Director, Principal Research Scientist, and Mike Titus, Staff System Administrator.

Q: How is your support organization structured?

We are an off-campus research station of MIT, employing about 100 people, with a mix of scientists, engineers, and programmers. The group is fairly loosely structured, with about 80 PCs and 10 workstations sharing a multi-segment ThinNet LAN. We employ one full-time system administrator who supports about 80 users and all workstations and network services at the site.

Q: What areas do you support?

Everything—installing and patching operating systems on workstations and servers, installing layered applications and public domain utilities, providing Help desk support for end-user questions, managing printer/plotter queues, performing backups, and providing hardware support. As for networking, we do everything from setting up the physical network to configuring and managing network resources and handling security for user data and the network.

Q: What particular challenges does a one-person support department face?

Our system administrator is on call essentially 24 hours a day. He gets many after-hour calls, which he can handle because he's usually at work even after hours. (Alan describes Mike's job as a "40-hour day.")

Balancing priorities can be difficult at times.

The Observatory is open around the clock, but since the work is scaled back in the evenings, that's when we do upgrades, network recabling, and other tasks that can disrupt operations.

We use scripts to automate administrative tasks. We don't do much remote management because all our computers are contained in one building; it's just as easy to walk over to a machine or to ping it.

We have a fair number of UNIX-savvy users, which at times helps to subtract from the workload, and at times does just the opposite!

Q: What kind of computers and operating systems do you support in your environment?

We have mainly HP Series 700 machines (two 715s, a 735, and a 712), plus a couple of older Series 800 machines (Model 835) and even an old Series 300 computer. We are running HP-UX Releases 9 and 10. In addition, we have four older (260) Sun machines and three SPARCstations running SunOS 4.x.

We also have 15 HP X terminals. We also have many Intel PCs running mostly DOS or Windows 3.1 and Windows 95 (with a few Linux, OS/2, Windows NT mixed in). We also have an HP C110 (our most powerful machine!), which is due to be updated to a C160 real soon now.

Our site started with HP 1000 machines in the mid-1970s, and a half-dozen are still in place for a number of critical real-time data-acquisition and control applications, but all software on these older systems is frozen and all new development efforts are targeted to HP-UX or HP-RT machines.

We also make heavy use of several HP 743s, running HP-RT, for controlling

real-time applications in conjunction with a suite of host HP 700s.

It's quite an eclectic collection, as pretty much every machine is tuned for its particular job, and we have to squeeze a lot out of relatively limited resources.

Q: Can you describe a typical desktop system or systems in your environment?

The typical desktop system is a PC or X terminal (depending on the nature of the work the user does). PCs run the gamut from 386/25 to Pentium 133s.

Engineers typically run the high-end PCs, the ones with greater power, memory and disk capacity; all have color monitors ranging from 14 inches to 21 inches, depending on usage. The scientists mainly use X terminals with moderate amounts of memory. Main disk storage is primarily provided by the 700 machines, with the heavy use of NFS cross-mounting. The PCs share files through these NFS servers as well.

No user has a full-fledged UNIX workstation at his desk solely for his personal use. Workstations are placed in public spaces and usually take on lots of "server" roles.

We have one HP 735 that provides many of the centralized services like disk space/printers/compute power, but other workstations provide some of these services as well. All of the workstations are stand-alone, and we don't centralize accounts through NIS, so all the workstations can function somewhat independently. Almost all "servers" are Series 700 workstations, although we still use Sun SPARCstations for some functions.

Q: Can you describe your printer environment?

We have group printers with multiport spoolers attached to multiple PCs, printers

Hewlett-Packard 9000

It's Our Specialty

200/300 Series	400 Series	700 Series	800 Series
216/236/217	425e	705/710	E/F/G
310/320/330	425t	715/720	H/I/K
350/360/370	425s	730/750	
318/319/340	433s	735/755	
345/375/380		C/J Series	

We also carry memory and interface for all of our workstations

Printers	Mass Storage	Plotters
2225A/B/C/D	9121/9122	DesignJets
3630A PaintJet	9153A/B/C	DraftPro
C1602A PaintJet XL	C2254HA	DraftMaster
LaserJet II/IIID/IIP	C2440HA/JA	Desktops
LaserJet III/IIID	C3232A	Electrostatic
LaserJet IIIIP,IIIsi		
LaserJet 4L/4P/4+/4+/4SI/4V/4MV	C2213A/D	
	DAT Drives	
	CD ROMs	
	Optical Drives	

Specials
DesignJet 650C Plotter



We offer large discounts, outstanding service and immediate delivery.

TED DASHER & ASSOCIATES

PH: 800-638-4833
FAX: 205-591-1108
E-mail: sales@dasher.com



CIRCLE 49 ON READER SERVICE CARD

RECERTIFIED HP EQUIPMENT

HP 9000

BUY ▼ TRADE ▼ RENT ▼ SELL

■ SERVERS	■ DISK DRIVES	■ MODEMS
■ WORKSTATIONS	■ CD-ROM	■ PCs
■ PLOTTERS	■ TAPE DRIVES	■ X TERMINALS
■ PRINTERS	■ MEMORY	■ UPGRADES

800-VIA-TECH

8 0 0 - 8 4 2 - 8 3 2 4



VIA TECHNOLOGY

603 NAVARRO, #800
SAN ANTONIO, TX 78205
PH: 210-227-7726
FAX: 210-227-6223



CIRCLE 127 ON READER SERVICE CARD

attached to workstations accessed by PCs and other workstations, and one high-speed, high-volume printer on the network used by all.

Q: What does your network look like?

The whole complex has ThinLan cable running 10MB/s Ethernet. We have one HP 4-port hub, which spreads four lines throughout various parts of the building. The most heavily populated section of our building has two separate lines running through it, and the traffic for each is isolated with HP bridges. We also link to a couple of other buildings through fiber-optic cables/repeaters.

We have a 10-port Livingston Communications Box that provides SLIP/PPP remote access. Our access to the Internet goes through a fiber-optic link to a Cisco router in another building where we share a microwave T-1 link with several other facilities for connection to the MIT campus (about 40 miles away) and the Internet.

Q: What communications services do you support between your users and between your users and their associates?

Most communication is done through e-mail locally and among associates worldwide. We use UNIX send-mail/SMTP exclusively (no PC mail gateways). We have lots of site aliases for users and groups, as most internal info is distributed via e-mail. We have an anonymous ftp site and a Web site for exchange of data and information with the outside world as well. We've been on the Internet for more than 10 years.

Q: Can you describe your backup strategy?

All file systems are backed up once per week, with the exception of a 20-GB Falcon RAID array, which is backed up only once per month. Nightly incrementals are run on file systems that

change often or are critical.

For backups, we use an HP 6-tape DAT autochanger and UNIX dump. Backups are mainly done to protect against loss of data in the event of disk failures.

Haystack Observatory is primarily a research facility, so there is generally no need for live online backup. Users must simply live with the inconvenience of having to restore data from backup tapes.

Q: How do you authenticate your users on your networked computers?

We use NIS on our older Sun machines, but on the HP machines we hand edit and copy password files in order to eliminate the need for a central account server and to better control access to each machine. We have a strict password policy and do password cracking with "crack."

Q: Can you describe how you manage your printer queue?

On the printers attached to UNIX boxes and the network printer, we use the UNIX lp spooler. The PC printing through the network is managed by PCNFS.

Q: What kind of training do you support for your users and do you yourself participate in?

There is no formal program. Typically people participate in classes or seminars whenever they relate to current tasks and/or on an "as needed" basis. Most learning is done through reading books and manuals.

Q: Can you describe the types of applications that are run in your environment?

The applications are pretty much evenly split between applications developed in-house for our particular scientific field and commercial packages for engineering/design work.

Our scientists develop their own software, and use the HP equipment in real-

time data acquisition and offline data analysis. Our engineers mostly use commercial packages, a wide variety of electrical/mechanical/CAD tools including Synopsis, WorkView, PowerView, and AutoCAD. We do a lot of custom VLSI work, using software developed at the University of New Mexico for ASICS integrated circuits.

We also make heavy use of freeware, particularly utilities.

Q: Can you describe the marketplace of your organization?

We are a scientific research facility. Visit our home page, <http://www.haystack.edu/haystack> for more information about our site. ■

Judy Tarutz is vice-president of Ibis Communications Inc., a technical documentation and training consulting firm in Lynnfield, Massachusetts. An internationally recognized expert in technical publishing, she is the author of Technical Editing: The Practical Guide for Editors and Writers (Addison-Wesley, 1992). Before forming Ibis, she worked as a Learning Products engineer at HP's Chelmsford facility and as a principal technical editor for Apollo Computer.

MINISOFT IS
TOTAL CONNECTIVITY



MIS

HP700/98, VT320, TN3270 & TN5250

NS/VT, TELNET & LAT

TCP/IP INCLUDED

ONLY \$169

COMPETITIVE UPGRADES JUST \$79

CIRCLE 74 ON READER SERVICE CARD

WINSOCK COMPLIANT • KEYBOARD REMAPPING • 32 BIT ARCHITECTURE

• DDE & OLE AUTOMATION • NETWORK FILE TRANSFER (FTP & NFT)

• RUNS ON DOS, MACINTOSH, WINDOWS, WINDOWS 95 & NT

800/682-0200

Internet address: sales@minisoft.com

MINISOFT

312 Maple Ave. Snohomish WA 98290

European sales: MiniSoft Marketing AG 41.41.340 23 20

MiniSoft 92 is a registered trademark of MiniSoft Inc. All other trademarks are the property of their respective owners.

What's a Domain?

I'D LIKE TO TALK ABOUT domain. Umm ... you mean that great operating system? No. Oh ... then you mean Internet domains? Sorry, I mean Microsoft Windows NT domains. Boy, is this confusing!

Well, there may be some confusion. While the name *domain* makes sense, it is a pity that Microsoft chose it—especially when NT machines are now so commonly connected to the Internet.

There are two ways a user can log onto NT: as a workgroup user, or as a domain user. A workgroup is useful if your machine is a stand-alone computer or if your work environment is a small group of loosely coupled users sharing resources. Workgroups are also useful if there are legacy Windows systems involved in the network environment. There is, of course, no central administration and each user is responsible for maintaining his personal machine and its environment.

A Windows NT domain is a unit of security. It provides user validation, centralized administration, resource control, resource sharing, and user environment configuration. Domains can be viewed as logical entities to which users subscribe for enterprise services.

This is why the NT log-on screen has three fields: log-on name, password, and domain name. If the domain name is the name of a local workgroup, the user is validated on the local workstation and limited to accessing only the local computer and those resources supplied by other workgroup members. When users log onto the domain, they have access to all of the resources of the domain. In fact, a domain log-on can access the resources of that domain, workgroups, and possibly the resources of other domains (We'll talk about this in a moment).

There must be at least one NT Server acting as the Primary Domain Controller (PDC) for an NT domain to exist. The PDC is the server where the user validation database is kept and which validates a log-on request to the domain. The log-on request asks the PDC to obtain the security descriptor and rights granted to a particular user. There are two other ways an NT Server can be used in the domain: as a Backup Domain Controller (BDC) or just as a server. The BDC is a server capable of serving as a PDC should the PDC become unavailable. The domain can have more than one BDC, but only one PDC. If the PDC fails or becomes unavailable to the network, the other BDCs have an "election" and one of the BDCs is promoted to Primary Domain Controller. This fail-over process happens so quickly that a user logging on won't even know it happened. And you don't have to worry about synchronizing any of the user information between the PDC and the BDCs because NT replicates any changes between them automatically.

Other than being configured as either the PDC or a BDC, the third way an NT Server can be used in the domain is as just a plain server. While enterprise services such as file servers, databases, and e-mail can be used on either the PDC or BDCs, it is sometimes wiser to place them on a server dedicated to the individual tasks, free of enterprise maintenance functions. Generally, dedicated servers are used when, for example, the database is heavily accessed. Not that PDC duties load a computer that much, but they can slow down users logging on and accessing other resources.

A note of caution! You cannot convert a server to a PDC or BDC, or a PDC or BDC to a server without reinstalling the

NT system. So plan your enterprise carefully to avoid redoing your installations.

I said before that NT domains allow the administration of user configuration environments. What I meant by that was two things. First, the administrator (the equivalent of a superuser) can define and restrict the user environment, defining what appears on the user desktop, what resources are available and to what extent they can be accessed. Second, all users can customize their desktop and environment and save it on the domain controllers. Your settings are then loaded onto whichever workstation you happen to log onto, giving you warm and fuzzies by supplying your favorite window colors, wallpaper, screen saver, and paths.

Before a user can log onto the domain from a workstation, that workstation must be a member of the domain. A workstation joins a domain either from a domain controller by having the administrator add the workstation computer name to the domain, or from the workstation by entering the domain name in the network options and rebooting. If a workstation is not in the domain's list of allowed workstations, then even a valid user cannot log onto the domain from that workstation.

There can be up to 40,000 users in a domain; a practical limit might be 10,000 to 15,000 per domain. That's large enough to house any company department, and many entire companies. Most small organizations tend to use one domain name for all users, but larger organizations separate the various departments into different domains, such as engineering, sales, and accounting.

Now, if you have an organization that is separated into different domains, wouldn't you like the various depart-

ments to communicate with each other and share resources? What about when people from different domains are working on the same projects? The solution to this communication and resource sharing among domains is what is called the Trust Relationship.

A Trust Relationship is a link between domains wherein one domain grants access to the users of another domain. The "trust" name comes from the fact that the first domain (the trusting domain) is trusting the second domain (the trusted domain) to have securely authenticated its users. The system administrator defines which domains his domain will trust. By default there are no Trusts. Each Trust is only one way. However, a Trust can be established in which two domains trust each other. This is called a two-way Trust and is composed of two one-way Trusts.

Using Trusts, you can create various configurations to manage your users and your enterprise's resources. Microsoft has grouped them into four basic configurations: Single domain model, Master domain model, Multiple Master domain model, and Complete Trust domain model.

The Single domain model is the model a small company might use, particularly if all employees are located at a single location. Resource sharing is not an issue because the domain contains all the resources as well as the users.

A Master domain model is composed of one master domain containing all the users, and one or more resource domains that contain the file and print resources. Centralizing the user database makes it easy to maintain the list of users, and grouping the resources into specific domains allows the resources to be managed by separate

groups within the company. Each resource domain trusts the master domain, which contains the user accounts.

Larger organizations might use a Multiple Master domain model. This model is similar to the Master domain model, except that there are more than one master domains. This becomes necessary when the total number of users exceeds 40,000 (the user limit for an NT Server). Or it might be more practical for companies whose offices are widely separated around the country or globe. In this model the users are contained in one of several master domains. The resources are contained in the resource domains as before, and each resource domain has a separate trust to each of the master domains.

The last model, the Complete Trust domain model, can be viewed as a collection of Single domain models coupled together by two-way trusts from each domain to each of the others. This model is for those companies organized into fiefdoms who cooperate grudgingly with each other. The potentially large number of trust relationships and lack of central administration make this one of the less desirable models.

I hope you now have a basis on which to plan your domains. In the next column, we'll talk about user accounts. ■

Bob Combs is the Director of Systems Architecture at PCSI in Englewood, New Jersey, a company specializing in client-server technology. He is a Microsoft Certified Systems Engineer (MCSE) and holds a master's degree in electrical engineering.

by Larry Headlund

The Future of X

BY THE TIME YOU READ this, the last X release, which was code named Broadway and which will probably be called X11R7, will be on the streets. The code name Broadway will probably have the correct resonance with you if you are intimately familiar with the geography of the area of Cambridge, Massachusetts, where MIT is located. Broadway is one of the main streets around MIT, the original home of the X project. Last time I spoke of what this release isn't—it isn't a requiem for X. This time I want to talk about what it is, as well as I can tell at this date.

Web, Web, Web

When I think of events I missed the significance of at the time, high on the list would be the talk on the World Wide Web concept given at the San Antonio Usenix conference in the early 1990s. This was one of the first presentations of the WWW outside of the Internet itself. I thought it was a mildly interesting way of programming in X, competing with all the other ways bubbling up at the time, such as TCL/Tk, Winterp (a Lisp interface), WCL, etc. I missed the significance of the document embedding paradigm—of how that concept would grab hold of the collective imagination. I certainly didn't think that a few years later "Home Page" would be part of the common discourse, that cartoons besides *Dilbert* would make references to it without explanation needed, that Republican presidential candidate Bob Dole would use some of his precious debate time to give out his home page address. In the last half decade WWW has conquered the world.

I should not have been so surprised. The WWW embodies a Big Idea. The Big Idea in computer science before it

that affected us all was the UNIX idea that everything is a file, everything is a data stream. The corollary to this is that the output of one file can be the input to another. Hence was made possible the chains of simple programs to do complex tasks. This in turn enabled the real use of the computer science principles developed in the sixties. UNIX was the first demonstration on a large system that these concepts could work, that they were not just academic exercises. It has been more than 25 years since that Big Idea, past time for another one. This is particularly true because the data stream idea did not translate well into the world of GUIs. The output of an X program is different in kind from its input. The Big Idea that WWW is about is the document-centric viewpoint—that it is the data, not the programs, that are central.

The document-centric approach was in the air in the early nineties. Certainly the data flow analysis and object oriented concepts predated it by decades. But we were seeing a lot of new expressions of this idea, not least because it seemed to marry so well with the notion of clicking on a document. The Standardized General Markup Language (SGML), which is the source of the HyperText Markup Language (HTML), was a standard to allow interchange of documents between word processors. The Hypertext idea, old in itself, was getting new power from environments—Mac, X, MS Windows—that could support it. The Open Doc and CORBA initiatives were brewing. Put this together with the Internet and you have the concept, if not the execution, of the WWW. In retrospect it seems an inevitable winner.

The WWW was designed as a tool for researchers, people who are oriented

towards text and data. Not all the data was textual, of course. Pictures, graphs, and formulas were always accommodated. But there certainly wasn't an emphasis on having snazzy-looking WWW graphics; the content was what was important. A popular way to access the WWW for the bandwidth-challenged was through Lynx, a text only browser.

At first, WWW was an X thing. Mosaic, sire of Netscape, was written in Motif and ran only with X. When the class of browsers expanded beyond "our" universe into the wider world, however, things really took off. Of course, a big spur to this development was the existence of an exciting, popular (in the sense of appealing to non-technical types), growing technology that was not controlled from the state of Washington. It inspired dreams of overthrowing Bill Gates or the even more tantalizing dream to some of becoming the next Bill Gates.

Browsers can do more than look at Web pages, but does anybody think they would be in use by millions if they were just a different way to read newsgroups and download files? They can be a paradigm to replace the desktop, but there wouldn't be the urge to stretch their capabilities without all those pages.

Rather quickly it was discovered that the current structure for browser use had some problems.

Isn't Java the Answer?

To virtually every computer problem associated with browsers, Java is offered as the answer. This is rather remarkable for a language that was designed for embedded systems, for running the proverbial toaster, but Java jumped on the Web wave with a vengeance. Do you need an object-oriented language? Java

is that. Do you find C++ too baroque? Java is similar but simplified. Too much bandwidth wasted sending images when a description of the image would be more efficient? Java can do that. Want to do something on the receiving end? Write a Java applet. You realize this could cause a security problem, particularly if you are sitting in front of an operating system with no real security? Java will have its own security. Someone has found a hole in Java security? That will be fixed. And so on.

Beyond the problems I have hinted at, Java has a big hurdle to face: it is new. That can be good if it means you can revisit old problems with new skills, but it does mean you have to revisit them. There is no one with long experience in Java, and too few with enough. And taking the Java path means you have to recode your application. You have to find or train Java programmers. You must keep training them, because Java is an evolving, rapidly evolving language. The code writing, debugging, documentation, and testing tools for Java were nonexistent a couple of years ago, though you can be sure this lack is being rapidly filled. Even with these new tools, you will have to buy them or upgrade your existing tools to a Java version, learn the new tools or new versions, and train on the tools as well as the language. Let me say again, you have to code to convert to Java.

X to the Rescue?

The curse of the computer world is legacy applications. It is a curse when there is movement in the air, anyway. I wouldn't want to have to re-create the universe every morning, but having to deal with the past when in the grip of a new idea is sobering. What is to be done?

The Broadway initiative says it has found a way out. Your existing X applications can be Web efficient, can interact with browsers without sucking up all the bandwidth, by linking to new X libraries. These libraries will be browser aware. When a browser is pointed at an X application linked with these new libraries, the application will send out HTML signals, not the familiar X signals. Note that the browser doesn't change, the browser doesn't have to be relinked—the application is conforming to the browser's expectations.

Every X application can be accessed from the browser. Existing X tools, X code, and X trained programs can be used without a change. This will get your order entry applications on the Web fast.

Now this does not do everything that Java does. You won't be running applets through the browser and hence a lot of the vision of a new model of computing isn't realized. However, it will save a lot of converting time and expense for existing applications. New applications that don't need the more revolutionary features of Java can probably be constructed faster this way. If nothing else, it may buy you some time to see how this revolution progresses before re-making your world. ■

Larry Headlund is the president of Mathematical Engineering, Inc., a UNIX and Motif development company and has been working with commercial UNIX since 1983 and with X since 1987. He can be reached at lmh@world.std.com or at (617) 242-7741.

by Joseph Berry

I HAD BEEN LOOKING for a way to generate graphs from within Perl. I had data I wanted to display through a Web browser and Perl was my CGI language of choice. The O'Reilly & Associates' book *CGI Programming on the World Wide Web* pointed me to two methods of generating graphs. The first was to use the GNU gnuplot program (I discussed this program in a previous column). With this technique, you pass all the required information to gnuplot, it draws the graph, and you display the output.

The book discussed a second method, using the package `pgperl`, which allows for direct Perl access to the building of the graph via procedure calls. This sounded easy enough: get `pgperl`, install it, try it out. That was when things started to become complicated. `Pgperl` isn't really the graphic library I was after. It is a Perl interface to a subroutine library called `PGPLOT`. `PGPLOT` is a well-known graphics library in the scientific community currently being maintained by Dr. Tim Pearson of the California Institute of Technology (tjp@astro.caltech.edu). It wasn't too hard to find the latest copy of this package (see below). Unfortunately, though, `PGPLOT` was written in FORTRAN. In this month's column, I discuss how I eventually got `PGPLOT` properly compiled and running.

A few months ago, I described a UNIX program called `xpaint`, a graphical editing tool. I subsequently received an e-mail letter from someone pointing out to me that there was a better program than `xpaint` that performed approximately the same functions.

Unfortunately, I somehow lost this person's e-mail so I don't remember who the individual was (although I believe he works for HP), or what the name of the program was. I think I

solved the latter problem of identifying the software by simply assuming there was such a thing and doing some snooping around until I found it. I believe the program that was recommended to me is called `xfig`. I discuss this program below. I would really appreciate hearing again from the individual who recommended the program to me. I want to thank him for his suggestion.

MISC

f2c

If you want an easy-to-use method of converting FORTRAN programs to C programs, look at `f2c`, a FORTRAN to C converter. This program is available by anonymous ftp from netlib.att.com. Look in the `/netlib/f2c` directory. The latest version of the program that I found is 19960619 (see `Versions.c`). Unfortunately, I had found an older version first and had tried it to convert `PGPLOT` to C. I was almost successful—near the end, `f2c` had difficulty with some esoteric FORTRAN 77 syntax, forcing me to look for another solution. I didn't come back to `f2c` until I had to search for this program's source for this column. It was then that I realized I never had the latest version. `F2c` even comes with a script file called `f77` that allows you to pretend you have a FORTRAN compiler where in reality the script converts your FORTRAN program to C and compiles it.

With this package, you will build and install two libraries, `/usr/lib/libf77.a` and `/usr/lib/libf77.a`, which consist of a number of support routines that the newly created C procedures use.

g77

After speaking with Dr. Pearson about my problems with `f2c`, he suggested I

ALPHANUMERIC PAGING FOR UNIX

ROBUST, RELIABLE, USER-FRIENDLY DELIVERY OF MESSAGES ANYTIME, ANYWHERE

- Email forwarded to pager automatically
- Pages can be generated from scripts and network monitoring programs
- GUI and command line interface
- Works with any paging service
- Automatic email confirmation, history logs and error reporting
- Client-server technology
- Works with digital and alphanumeric pagers



**Personal Productivity Tools
for the Unix Desktop**

14141 Miranda Rd
Los Altos Hills, CA 94022
Email: sales@ppt.com
Tel: (415) 917-7000
Fax: (415) 917-7010
<http://www.ppt.com>

try getting g77, the GNU FORTRAN compiler. Was I ever embarrassed. I've been a GNU supporter for years but never knew that the Free Software Foundation had developed a FORTRAN compiler. This compiler is based on the GNU gcc compiler and you must have the source tree to the compiler available when building g77.

You can download g77 from *gatekeeper.dec.com* or any other site that maintains GNU sources. My copy of g77 was called *g77-0.5.18.tar.gz*. Follow the instructions carefully. The actual build takes place in the gcc source tree (don't use a version of gcc before 2.7.2).

PGPLOT

Before you build PGPLOT, you must decide what output displays you will use. Examples include Canon laser printers, HP plotters, HP printers, PostScript printers, X window dump files, Tektronix terminals, etc. Once specified, PGPLOT alters the Makefile to build the appropriate display drivers with the rest of the package. Before you proceed to the make, look at the files in directory *./sys_hp* as there are instructions and hints for how to build the software on HP systems using the HP FORTRAN compiler.

If g77 has been built and installed, PGPLOT follows without a hitch. Compiling the software with g77 was remarkably easy.

Except for one of the dozen or so test programs, everything compiled and executed without a flaw (the one exception had a few compile errors which were actually legitimate). Pretty graphs appeared on my X window representing mathematical functions, plotted data points, and X-ray sources from outer space.

Included with the software is a package

called *cpg*. This is a library that allows C programs to talk to PGPLOT. Be sure to build and install this package, too.

PGPerl

As I mentioned above, *pgperl* is a Perl interface library to the PGPLOT graphics library. *pgperl* is available for either Perl 4.0 or Perl 5.0. The software can be obtained from *ftp.ast.cam.ac.uk* in directory */pub/kgb/pgperl* or its mirror site in the U.S. at *linux.nrao.edu* in */pub/packages/pgperl*.

Once I had PGPLOT functioning properly, building *pgperl* was not difficult. You can build the software with dynamic libraries or static libraries. A suite of test programs comes with the package, showing you exactly how to use each of the PGPLOT features.

Documentation consists of a 279-line doc file.

The biggest fault I found with the *pgperl*/PGPLOT package is the level of

sophistication of the graphics functions. All the primitive functions are in place but there are not a large number of higher level functions. For example, you cannot draw pie charts (you can draw circles, arcs, straight lines, however, in order to implement a pie chart).

tcp_wrapper v 7.4

Today, anyone connected to the Internet needs to be concerned about security and access control. Do you know who is using telnet into your system or who is accessing your files via ftp?

Thanks to Wietse Venema (*wietse@wzv.win.tue.nl*) of the Department of Mathematics and Computing Science at the Eindhoven University of Technology in The Netherlands, we have a program called *tcp_wrapper* to help get a handle on remote access to your systems. *Tcp_wrapper* adds a small wrapper program around the common network services: finger, ftp, telnet,

CIRCLE 91 ON READER SERVICE CARD

rlogin, rsh, etc., etc. When a network service request arrives, the tcpd wrapper program intercepts the request, optionally logs it, and determines whether to permit the connection based on the contents of files *hosts.deny* and *hosts.allow*.

At my own site, I ordinarily disallow all access from "foreign" computers. I then permit access from our corporate headquarters but make tcpd send me an e-mail notification when a network service is invoked. It gives me a warm, fuzzy feeling.

You can get a copy of tcp_wrapper from info.cert.org. The file is called `/pub/tools/tcp_wrappers/tcp_wrappers_7.4.tar.gz`. Everyone should be using this program.

xfig v 3.1.4

Xfig is available from ftp.x.org as `xfig.3.1.4.tar.gz` in directory `/contrib/applications/drawing_tools/xfig`. The program is a menu-driven tool that allows the user to draw and manipulate objects interactively in an X window. Many example figures are included that demonstrate the power of the program. Documentation includes a 35-page man page document as well as lists of compatible tools the user may want to integrate with xfig. The program works with many of the standard file formats, such as gif. However, it uses a proprietary "fig" format as its main file storage repository.

COMP.SYS.HP.HPUX

xmkmf

On another note, someone had sent me an e-mail message asking how to get a copy of xmkmf for his HP-UX system. He was trying to compile some X applications that came with an Imakefile file.

Jeff Hoxworth from the HP Roseville

Response Center was kind enough to post a response to someone else's inquiry on the comp.sys.hp.hpux newsgroup. As he said, "If you don't have the programming environment, you can grab them from <http://hpux.cae.wisc.edu/>."

To get xmkmf, go to this site and choose "Catalogue Of Public Domain HP-UX Ported Software." Then choose "X11/Core." Download the imake-5.06 package (in source or binary) This package includes xmkmf.

COMP.UNIX.SOLARIS

FileRunner v 2.0

If you love the DOS/Windows Norton Utilities, FileRunner is for you. With one graphical/GUI front-end, you get commands like ls, mv, rm, cp. You can also browse ftp directories as easily as local directories. Other features include a recursive hyper-menu of the entire file directory tree for quick directory traversing, a hotlist of often-visited locations, and built-in command shell windows with filename completion.

FileRunner uses Tcl/Tk 7.5/4.1 and needs an ANSI C compiler for compilations. The software was written by Henrik Harmsen (hch@cd.chalmers.se) and is available from his Web page at <http://www.cd.chalmers.se/~hch/filerunner.html>. Via ftp, the software is also available from sunsite.unc.edu as `/pub/Linux/X11/xutils/managers/FileRunner_2.0.tar.gz`.

WEB PAGES

<http://www.omg.org/corbask.htm>

If you are interested in information about CORBA, the Common Object Request Broker Architecture, you can download a copy of the CORBA 2.0 specification by accessing this Web page.

<http://www.molink.com/~sclink/space>

SpaceBound is an online magazine dedicated to astronomy and space exploration (one of my favorite non-computer subjects). The Web site features software reviews, Internet links, news, articles, and much more. Be prepared for lots of graphics.

<http://www.boutell.com/faq/htedit.htm>

If you are designing Web pages, then you'll probably need a good html editor. This Web site lists loads of html editors for developing your Web pages on a PC, Mac, or UNIX box. Commercial and free software is listed.

<http://pms.com/cliff>

Yes, another mystery Web page. As the original announcement stated, "Cliff is a stylistically and structurally original online murder mystery/soap-opera/comedy which takes place in a small college town where a 25-year-old murder has been discovered. It's an enormous storyscape spanning decades and you'll need to cruise both space and time and meet dozens of characters to find out what's going on."


Joseph Berry is a senior software developer at Landmark Systems Corporation in Vienna, Virginia. He is one of the authors of Landmark's PerformanceWorks products, PerformanceWorks/Smart Agents for UNIX. A former HP 3000 systems specialist for Hewlett-Packard, he has been in the computer industry for more than 25 years. He can be reached at joe@wayne.unix.landmark.com.

Thinking about DLT?

Think again! We know that many of you are thinking about DLT tape libraries. But recent advances in 8mm technology and a library made by StraightLine—might make you want to think again.

	StraightLine	StraightLine	ATL	ADIC
Model	SL-850	SL-8150	ACL 4/52	Scalar
Drive format	8mm*	8mm*	DLT 4000	DLT 4000
Number of tapes in library	50	150	52	52
Tape drawer(s)	Yes	Yes	No	No
Removable tape boxes	Yes	Yes	No	No
Sustained transfer rate (native)	3MB/sec	3MB/sec	1.5MB/sec	1.5MB/sec
Total capacity (native)	1.25TB	3.75TB	1.04TB	1.04TB
Bar code reader	Yes	Yes	Yes/6-digit	Yes
Number of drives	4 or 5	4 or 5	4	4
Head life (hours)	>20,000**	>20,000**	10,000	10,000
Media uses (passes)	20,000	20,000	15,000	15,000
Drive MTBF (hours)	200,000	200,000	80,000	80,000
Library list price	\$43,995	\$61,995	\$65,000	\$65,000

* Exabyte's recommended 8mm tape—Exatape 170M Advanced Metal Evaporated tape
 Sony's recommended 8mm tape—SDX-13N 170M Advanced Metal Evaporated tape
 ** Exabyte Mammoth head life—20,000 hours
 Sony AIT SDX-300 head life—30,000 hours



Led by Exabyte and Sony, 8mm technology has taken a giant step forward in performance and capacity—3MB per second data transfer rate and up to 25GB per tape in native mode. With built-in compression, it's even better. And that's not all.

Reliability — DLT, once the leader in reliability, has been surpassed here, too. StraightLine took the best in 8mm drive technology and engineered a library with mainframe-class robotics— aerospace structural frame and optical sensing of mechanical movement. It's a powerful combination of reliability and performance at a really competitive price.

Easy to maintain — It can be expensive to replace DLT library drives. They have special tape ejection motors and a modified handle, and you can only get them from the library manufacturer. If you ever need to replace a drive in a StraightLine library, just remove four wing nuts, unplug the drive, and pop in an off-the-shelf replacement.

Built with the operator in mind — With seven years of experience manufacturing tape libraries, StraightLine knows how to make one that's operator friendly. Unique tape drawers slide tapes out of the machine to the user. There's no reaching in and around the robotics. It takes only seconds to exchange one tape or all the tapes, which are loaded in convenient tape boxes.

Industry standards — StraightLine libraries are compatible with all standard systems and software—Sun, HP-UX, IBM RS/6000, Sequent, NT, DEC; Novell, LAN Manager; Arcada, Legato, ArcServe and many others.

If you're thinking about a tape library, think about 8mm. Then call StraightLine—the leader in 8mm tape libraries. Now that's good thinking!

Call: 1.800.458.1273



straightline

677 120th Ave. NE, Suite 146 • Bellevue, Washington 98005
 www.igm.com • e-mail: igm@igm.com

4mm format also available.

© 1996 StraightLine. All rights reserved. StraightLine is a division of IGM Communications.
 All brand and product names are property of their respective holders.

CIRCLE 165 ON READER SERVICE CARD



HP 1000 Guru

Q: I am transferring an ASCII file from my RTE-A system using ftp, to a PC running NT. The file on the 1000 has records that are greater than 256 words (512 bytes). I find that the records are being truncated at 512 bytes. How do I override this limit?

A: When transferring ASCII files using ftp, the default record length is 256 words (512 bytes). If a type 3 or 4 file on the 1000 has a record length greater than 256 words, transferring this file to a non-HP 1000 system will truncate the record length to 256 words.

To override the 256-word default when transferring an HP 1000 file to a non-HP 1000 system, you must specify the full name of the source file. For example, suppose you have a file on the 1000 that looks like:

			siz	blk	words	recs	rlen
LONG.AS	rw/r	4	24	24	402	1	400

To transfer the file using ftp on the 1000:

```
ftp> put long.as:::4:::400 long.as
```

One would have thought that ftp would default to the attributes of the source file, but this is not the case. One must specify the attributes of the source to avoid record truncation.

Q: I have increased the number of TELNET LUs in my system to 30. I now find that when my system gets busy, users are unable to log on. I have increased the number of TELNET LUs that HPMDM can handle. What else should I be looking for?

A: NS/1000 has a configurable limit on the number of User programs that may be active at any one time. This limit is typically around 25. Thus if 25 TELNET LUs are concurrently connected, no more will be allowed until one or more terminate. This condition can be detected two ways:

1. The NS Event log file will record an error that looks like this:

```
INETD RsrceLim SIGMOD
```

2. The INETD log file, */etc/inetd.log*, will contain errors:

```
NS error 122(4) scheduling TNS.T.
```

If we look up NS (NetIPC) error 122(4), we find this is "Too many users." The number of user records is configured by NSINIT and can be adjusted accordingly.

Q: I have recently updated to RTE 6.2 and NS/1000. I also use the 12079A Direct

Driver access for a LAN application written years ago. The problem I am now seeing is that the 1000 LAN no longer is able to receive a broadcast packet. Directly addressed packets still work. If I shut down NS/1000, then my application works. What changed at 6.2 to cause this problem?

A: What changed was that a bug in the LAN driver, ID*67, was fixed, which then brought to light a bug in the LAN firmware.

NS/1000 enables ARP Packet Filter Mode by issuing the CN command to the driver. The driver then checks the revision of the firmware on the LANIC. If the revision of the LANIC firmware is = or > than 6100, the firmware is instructed to enable ARP Packet Filter Mode. If the firmware is pre-6100, then the driver takes care of the ARP packet filtering.

The problem in the LANIC firmware is that when ARP packet filtering is enabled, it also blocks broadcast packets.

Now for the reason 6.1 works: ID*67 (the LANIC Driver), when checking the revision of the LANIC firmware, would check to see only if the firmware was = 6100, as opposed to = or > than 6100. Thus with 6110 firmware, the revision test would say the card was old and the ARP packet filtering would be performed by the driver—which works.

There are two workarounds:

1. Disable ARP Packet Filtering via the CN command:

```
CN, <LANIC_LU>, 36B, 0, 0
```

2. Install 2547 (old) LANIC firmware.

Solution #1 means the system would *not* have any ARP Filtering, and thus performance could be affected.

Solution #2 means the ARP Filtering could be enabled, and the driver would do it. But it would cause support issues since the old 2547 firmware can no longer be ordered.

And, of course, solution #3 would be to fix the firmware.

And solution #4 would be a special modified ID*67 driver that would always do the ARP filtering itself, instead of in the firmware.

This is a new problem, and as of this writing, no factory solution has yet been determined.

Q: I am adding a new SCSI disk to my RTE-A system. I used the total usable blocks as specified in the manual that came with the disk, but that number is apparently incorrect. Is there an easier way than trial and error to determine the size of a particular SCSI disk?

A: Use VSCSI. Assuming you already have at least one SCSI disk LU genned, use the *UNITSIZE* command as follows:

```
CI> VSCSI 25
```

```
VSCSI: UNITSIZE
```

```
Information on SCSI disk UNIT containing LU 25:
```

```
-----  
Number of PHYSICAL blocks on this unit: 2647080.
```

```
Size of each PHYSICAL block on this unit: 512 bytes.
```

```
This is equivalent to 5294160 RTE blocks ( 256 bytes per block)  
-----
```

Then use the number of RTE blocks, 5294160 in this example, for creating your disk LUs.

Q: Could you provide an updated list of the CI REVLEVEL values? The last list I have only goes up to 5270.

A: Sure. Here is a list starting with 4010 through 6200:

CI REVLEVEL	OS Revision
<851022.1015>	4010
<870305.1535>	5000
<881004.1257>	5010
<900306.1214>	5020
<911121.1717>	5270
<921124.1404>	6000
<930804.1632>	6100
<950206.0858>	6200

Q: Great! Now, can you do the same for the timestamps of BOOTEX?

A: Here we go:

BOOTEX	Timestamp	Revision	Notes:
11:21 PM TUES	4 FEB 1986	2540/4.0	Mux at select code 20 works
5:36 PM MON	3 NOV 1986	4010	Bootex grows to 768 blocks for D mux and Datapair support
7:57 AM THU	1 OCT 1987	5000	ID segment change from 4010
7:24 PM WED	30 NOV 1988	5010	Speed sense disable works
1:23 AM SAT	21 APR 1990	5020	Support for console-less systems
1:51 PM FRI	23 AUG 1991	5250	First SCSI-compatible bootex
1:19 AM WED	8 JAN 1992	5270	SCSI-compatible bootex shipped with RTE-A
1:46 AM WED	2 DEC 1992	6000	6.0 ID segment change
1:46 AM THU	25 NOV 1993	6100	Booting Datapair can take long time if disk is down.
1:39 AM THU	4 MAY 1995	6200	LDR ERR 411 with non-zero SCSI passthrough fence fixed.

Q: When is the next release of RTE-A expected?

A: The next planned release is due early in 1997 and will be a patch release only. It will incorporate patches to existing problems.

Q: Since upgrading to SCSI disks, I have been somewhat dismayed to see about the same performance as with my old CS/80 HP-IB disks. Comparing the specifications between the two leads one to expect the SCSI disk to outperform the CS/80 HP-IB disks. Is there a reason why the performance is not greater, and can anything be done to fine-tune the SCSI disk performance?

A: One of the major differences between the newer SCSI disks and the older CS/80 disks is the physical sector size. From at least the 7900A disk drives of the early 1970s until recently, the physical sector size was 128 words (256 bytes). RTE relies heavily on this "constant," and to modify it would require extensive rework. Thus, when the SCSI project was undertaken, one of the major concerns was to make the larger sector sizes (512 bytes for the current hard drives and 1024 for the MO drives) transparent to the OS. To accomplish this, the SCSI card uses an on-board cache on a per channel basis. The management of this cache can either increase or decrease performance, depending on how the disk is utilized.

Decreased performance due to the different sector sizes can be demonstrated by looking at the following comparison of the blocks:

256 bytes	256 bytes	256 bytes	256 bytes	...
◀— SCSI Block 0 —▶	◀— SCSI Block 1 —▶	—	—	...
RTE Block 0	RTE Block 1	RTE Block 2	RTE Block 3	...

Comparison of a 512-byte SCSI block with a 256-byte RTE block

When an RTE user wants to write to RTE Block 1, the following actions must occur:

1. SCSI block 0 is read
2. Part of the block is modified
3. The block is written back to disk

When an RTE user wants to write 512 bytes to RTE blocks 1 and 2, the overhead becomes even more complex because of the activities required to read 2 SCSI blocks (1024 bytes total), modify the needed portions, and write the blocks back to disk.

This problem can be avoided if all transfers are aligned so that they begin and end at SCSI block boundaries. Files can be forced to start at SCSI block boundaries by genning disk LUs so that the allocation unit times 256 bytes is equal to the SCSI block size. The current hard disks (C2212A and C2213A) have 512-byte blocks. This means an allocation unit of 2 will force all files on the disk to start at a SCSI block boundary. The allocation unit is the number of blocks represented by each bit in the bitmap at the start of each CI volume. (The allocation unit is further described in Chapter 10 of the *RTE-A System Design Manual*, p/n. 92077-90013.) The bitmap has a fixed size of 8192 words (128k bits), and each bit specifies whether each allocation unit is used or free. A disk containing between 128k blocks and 256k blocks will have an allocation unit of 2. Note that in addition to having an allocation unit of 2 or more, the SCSI disk LU must also start at a SCSI block boundary to be aligned.

Tests comparing SCSI disk transfers with the equivalent HP-IB disk show that the HP-IB disk will always outperform the SCSI drive when the RTE block is not

Are You Moving?

Send your change of address to:

Member Services

Interex

P.O. Box 3439, Sunnyvale, CA

94088-3439, USA

Or call:

408.747.0227

Or fax:

408.747.0947

Or e-mail:

membership@interex.org

interex

Interex Online

<http://www.interex.org/>

The International Association of Hewlett-Packard Computing Professionals

NEXS Spreadsheet

FAST 100% UNIX;
no emulators!

FLEXIBLE Customizable
functions and script-language.

conNexions API Let's you plug in
applications locally or over the 'net.

INTERNET-READY 63 links for sharing
data with other applications.

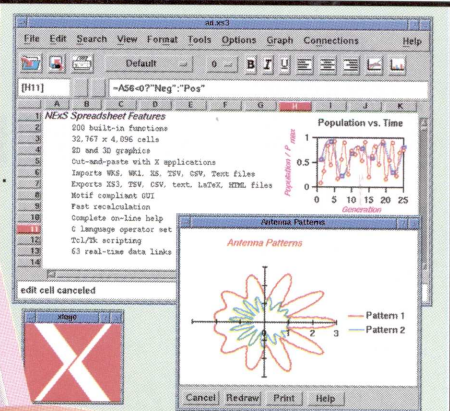
POWERFUL 237 built-in functions.

GRAPHICS 13 types of 2D
and 3D graphs.

IMPORT/EXPORT Excel and Lotus files,
Postscript, Netscape tables.

LOW PRICE Personal Edition - \$149
Professional Edition - \$249

PHN: 800-961-7840
INTL: 919-387-0076
FAX: 919-387-1302



NEXS

MULTI-PLATFORM
HP/UX, Solaris, SunOS,
AIX, Linux, OSF/1.

FREE NEXS Plug-Ins
at www.xess.com

FREE NEXS Demo
at www.xess.com

XESS Corp.

CIRCLE 132 ON READER SERVICE CARD

aligned with the SCSI block. This penalty can become significant when buffer sizes are increased. As a result, in the following discussions, when comparing the performance of HP-IB disks with SCSI disks, we will always be referring to transfers that have RTE blocks aligned with SCSI blocks. Also, all transfers are an even number of RTE blocks. This ensures that both the beginning RTE block and ending RTE block are aligned with the SCSI block. Note that for SCSI block sizes larger than 512 bytes, a larger transfer length will be required to maintain SCSI block alignment at both the starting and ending transfer points.

The effect that increasing buffer sizes has on performance depends on the type of file being transferred. For type 1 files, the HP-IB and SCSI disks have comparable transfer rates until the reads or writes are greater than 16k bytes. For HP-IB drives, buffer sizes >16k do not show an appreciable performance difference, whereas for SCSI drives buffer sizes up to 28k bytes continue to show performance increases. SCSI performance degrades when buffer sizes are increased beyond 28k because of an extra seek. Given the size of cache and additional buffer management, performance will decrease for SCSI reads when block sizes are between 6k and 10k and for writes when block sizes are between 8k and 10k.

For type 2 and greater files, disk performance is generally only slightly better for SCSI disks than for HP-IB disks when

transfer sizes are greater than 10k. However, for some transfer sizes of less than 10k, HP-IB drives may perform significantly better than SCSI drives.

In summary, when using SCSI drives, it is recommended that you use transfers that are SCSI block aligned. Starting block location can be forced to be SCSI block aligned by generating the disk larger than 128k blocks. Ending block can be forced to be SCSI aligned by specifying an appropriate transfer length. In general, performance will increase as transfer sizes (buffer sizes) increase. However, transfer size does not affect performance as much as SCSI block alignment does. In addition, when transfers are SCSI block aligned, performance differences between HP-IB disks and SCSI disks are only slight unless the file is a type 1 file and the transfer sizes are between 16k and 28k. For user applications, we recommend that individuals do performance modeling in order to determine which transfer lengths are most appropriate for their applications. ■

Walt Boeninger works in the HP Response Center in Mountain View, California. He has been supporting the HP 1000 for 15 years. His e-mail address is: walt@hpwrce.mayfield.hp.com



CSL Perspective

OVER THE LAST SEVERAL years, I have been writing and talking about the enormous change that affects us almost daily and have suggested some ways that we might adapt. A rapidly changing environment is difficult to deal with and even harder to keep track of—you want to throw up your hands in surrender. Against this stark reality, I've had one of those experiences that really bring home the need for continually expanding my skills and capabilities. Let me digress a moment and tell you the story.

I support some systems that could be termed mission critical. I'm on call around the clock on some of these environments and need to react to problems very rapidly. So I have one of those electronic leashes (a pager) and also a laptop so that I can dial in from almost anywhere. As luck would have it, my laptop had to go into the shop for repair, which means I had to fall back on my Radio Shack Color Computer, vintage 1985 and use a vt100 terminal emulator and 9600-baud modem.

So here I am, taking a day of vacation and siding my house, when one of our application consultants calls me to ask if I can postpone the backup from Friday night to Sunday. "Sure, why not," I said, before I realized that our backup tool uses a Motif/X11 user interface, which would be useless from home. Now the little wheels begin to turn as I'm trying to come up with a way to postpone the backup without having to drive into the office. I could remove the crontab entry that triggers the backup; that should be fairly easy. But wait just one minute! I wonder where the backup utility keeps the schedule information? After scanning a few man pages I discover that not only are the schedules in flat ASCII files, the format of the

schedule is very straightforward. So a few quick commands with vi, and mission accomplished!

As I went back to the siding job, I reflected a bit on what had just transpired. What skills did I use, what was my attitude and did it influence the outcome? Would I have handled it differently given a different set of circumstances? Would another system administrator have been able to pull it off? There are four key skills and attitudes that seem to come into play in this situation.

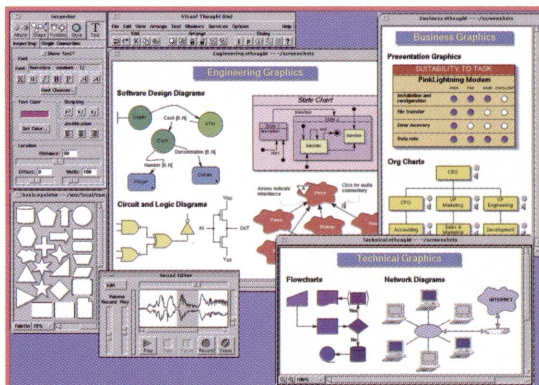
First and foremost is a "can do" attitude. I will not say to the user, "I'm sorry, it can't be done." I was positive that a way could be found, even if a drive to the office was required. This attitude is crucial in dealing with users, who are interested only in results and outcomes, not in the process or details of solving the problem. The "can-do" attitude often can signify a level of commitment to the users, to which they will reciprocate with a level of trust in your commitment to them. This relationship is what "customer focus" means to me.

Creativity plays a part here as well. Before one can be creative in solving a problem, one must be motivated to take a little risk with the creative approach. I've known some capable system managers who, when confronted with complex problems, throw up their hands in frustration because they are afraid to think creatively. They are more comfortable with what is obvious and certain. Sometimes certainty and stability will work very well, but it won't be enough.

Third, there's what I call "dogged persistence." Searching for answers to what seems an impossibility can be the biggest challenge. Plodding up a 12,000 foot mountain at age 38 seemed pretty

visual thought®

UNIX Diagramming & Flowcharting



Use Visual Thought for:

- Software design diagrams (Booch, Rumbaugh, Objectory, Fusion, custom notations)
- Flowcharts for ISO 9000, TQM, BPR
- Clickable Web diagrams with GIF/JPEG and server/client-side imagemap export
- Framemaker documentation graphics w/ MIF, EPSI export
- Network diagrams with network component clip art
- Dataflow diagrams, org charts
- Presentation & business graphics
- Graphical pre/postprocessor for other tools or simulators

FREE CD-ROM, FTP & 30-day trial

Additional benefits

- Intelligent, rubberbanding lines
- Macintosh-like ease-of-use... on UNIX!
- Editable, WYSIWYG drag-and-drop palettes
- Hyperlinked, hierarchical documents
- Arbitrary object rotation
- Complete text handling: subscripts, arbitrary fonts, sizes, colors, styles and justification
- Dozens of export formats: GIF89, JPEG, MIF, TIFF, EPSI, XWD, SunRaster, others
- Available on SunOS, Solaris, HP-UX, and soon, Windows 95/NT

VISIT (use code AE14)
<http://www.confluent.com/>
 800-780-2838 ext.154
info@confluent.com



Confluent, Inc.
 132 Encine Court
 San Francisco,
 California 94127
 415-586-8700 Voice
 415-586-8838 Fax

See why Visual Thought has thousands of users worldwide! Clients include:

ABB • AT&T • Alcatel • BNR • Bellcore • Boeing • CS First Boston • Credit Suisse
 E-Systems • EDS • Ericsson • Ford • HP • IONA • Lucent • McDonnell Douglas
 Motorola • Nokia • Nomura • Nortel • Persistence • P. Prediction
 Pure Software • SAIC • Smith Barney • Sprint • Toshiba
 U.S. Air Force • U.S. Navy • U.S. West • Wells Fargo • Nikko

©1998 Confluent, Inc. All rights reserved. Visual Thought, the Visual Thought logo, Confluent, and the Confluent logo are trademarks or registered trademarks of Confluent, Inc. All other names are the properties of their respective holders.

CIRCLE 37 ON READER SERVICE CARD

Do You Know Where Your Security Holes Are? Find Them with SecurityAudit/UX!

Have You Heard that UNIX is Notorious for Its Lack of Security Features?
Do You Know Where to Check to See if Your HP-UX System is Secure?
Do You Have the Time to Do This Checking Regularly?
Use SecurityAudit/UX To Do It All!

EVEN IF YOUR SYSTEM IS SET UP CORRECTLY (AND HOW WOULD YOU KNOW IF IT WAS?) AND HAS NO SECURITY LOOPHOLES, IT CAN BE VERY DIFFICULT TO MONITOR SYSTEM CHANGES, AND TO ENSURE THAT SECURITY ISN'T COMPROMISED. THE MAGNITUDE OF THE PROBLEM INCREASES AS THE TOTAL NUMBER OF USERS CONFIGURED AND THE TOTAL NUMBER OF FILES GROWS. IT'S EASY FOR ORDINARY USERS TO CHANGE THE SECURITY OF THEIR OWN FILES TO ALLOW OTHERS TO ACCESS THE CONTENTS. A LOOPHOLE LEFT BEHIND INADVERTENTLY OR ON PURPOSE MAY BE EXPLOITED BY A DISGRUNTLED EMPLOYEE OR A HACKER TO BREAK SYSTEM SECURITY, SOMETIMES MUCH LATER.

SecurityAudit/UX PRODUCES OVER 40 REPORTS, CONTAINING DETAILED INFORMATION ON THE FOLLOWING CLASSES OF PROBLEMS:

- User and Group-related problems, including weak passwords and non-unique identification numbers.
- File-system related problems, including historical tracking of files and detection of potential Trojan horses.
- PDF-related security problems, extended to detect changes in ACL specifications.
- Logging subsystems status display, and logfile analysis.
- Network-related status display and configuration weaknesses.

SecurityAudit/UX RUNS ON ALL HP-UX BASED 9000 SERIES 700 AND 800 SYSTEMS, AND HAS BEEN SPECIFICALLY TAILORED TO ADDRESS PECULIARITIES OF HP-UX, SUCH AS PDF, ACL AND HP'S SHADOW PASSWORDS.

Call EUGENE VOLOKH for more info!

VE*SOFT
 Purveyors of Fine HP Software
 Since 1980
 (310) 282-0420

1135 S. Beverly Drive
 Los Angeles, CA 90035 U.S.A.
 FAX (310) 785-9566

CIRCLE 40 ON READER SERVICE CARD

impossible to me. It was that persistent drive to beat that mountain that got me to the top. I didn't sprint, unlike some of the teenagers who were with me, but with a little more time and frequent rests, I got there. The great thing about it was the feeling at the top. Patience and motivation frequently go hand in hand with persistence. You can't give up, despite the frustrations and setbacks.

Finally, confidence in yourself, in your skills and capabilities, is what binds it all together. Unfortunately, this seems to come only with experience—having the opportunity to develop new skills, try things out, fail a few times and start again. As we grow in capability and skill, we naturally learn about ourselves and can begin to see clearly not only what we *can* do, but also what we *like* to do. That's when work begins to be fun!

A quick update on the CSL. As I've mentioned in previous columns, yours truly is working on the 10.0 porting efforts along with several other people. I've been working on some of the utilities we use internally with pretty good success. Our target date is to have around 100 of the most popular packages from the net available in time for the April '97 InterWorks Conference. This will coincide nicely with the release of the HP-UX CSL as well. Stay tuned to next month's column for a more detailed update on what's happening and my impressions of porting to 10.x. ■

Paul Gerwitz is chairman of the CSL committee and is a technology specialist at Eastman Kodak Company in Rochester, New York. He can be reached at 716-477-3067 or e-mail at gerwitz@interex.org or gerwitz@kodak.com.

Continued From Page 16

Enterprise Fax Management

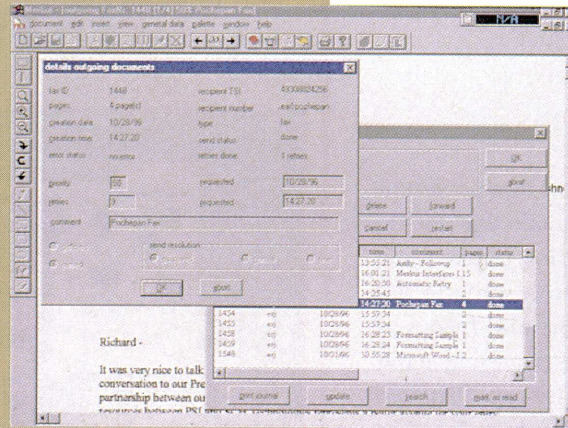
SCH Technologies has announced Merkur, a fax management product developed by Com-EM-Tex of Germany. SCH Technologies has exclusive U.S. marketing rights for the product.

Merkur features an exclusive Distributed Server Architecture that allows users to install fax resources such as modem servers, document servers, and print servers as modules. This provides unlimited scalability throughout the worldwide enterprise. The architecture also provides for distributing fax services around the enterprise to take advantage of available system resources and ensure full-time availability.

Running as a background process on available system resources, Merkur automatically and intelligently routes incoming faxes. Merkur's routing rules can recognize the incoming phone line (via Direct Inward Dial support), the document's source, and built-in bar codes. Merkur also provides automated outbound fax services including mass fax and faxback.

Currently running on UNIX-based servers, Merkur will be available for NT later this year. The price for one fax server and five users is \$1,500.

Contact SCH Technologies, phone: (513) 570-0455, fax: (513) 579-1064, e-mail: info@sch.com.



SCH Technologies, Merkur

classes available; Thread<Toolkit> for creating and synchronizing threads; Time <Toolkit> featuring 64-bit time and date classes; Network<Tool-

Kit>, a complete set of cross-platform classes for network socket programming; Web<Toolkit>, for creating sophisticated Web pages quickly and easily; and other libraries.

Contact ObjectSpace, phone: (214) 934-2496, fax: (214) 663-9100, <http://www.objectspace.com>.

Performance Management

PLATINUM technology, inc. has announced PLATINUM TransTracker and PLATINUM WireTap. TransTracker is a distributed systems transaction measurement tool that allows application developers to gather system, database, and network performance statistics throughout the development cycle. The product can also be used to evaluate the performance of third-party applications (such as financial or manufacturing applications) before an organization purchases them. Infrastructure architects use TransTracker to quantify the additional processing requirements and network traffic needed to operate a Java-based Web site.

WireTap captures and analyzes network traffic, providing a complete view of overall network resource usage and a detailed view of response times for individual application components and transactions. WireTap can identify the users visiting a corporate Web page and identify the sites

being visited by employees using a corporate server for Internet/Web access.

Contact PLATINUM, phone: (800) 442-6861 or (708) 620-5000, e-mail: info@platinum.com.

C++ Class Libraries

ObjectSpace, Inc. has announced the second generation of its C++ Component Series, which consists of 10 C++ class libraries, including all new versions of Systems<Toolkit> and STL<Toolkit>.

The libraries provide a full set of features from a portable implementation of the Standard C++ Library to libraries for systems programming and Web development. All libraries are portable across popular hardware platforms, operating systems, and compilers and are safe for use in multithreaded environments.

The library includes Standards <Toolkit>, said to be the most portable, high-performance implementation of the Standard Template Library (STL), string, standard exceptions, and utility

PPP Software

Progressive Systems Inc., the new home of Morning Star PPP, has announced the continued development, sales, and support of the software with a new release. Morning Star PPP is a point-to-point protocol (PPP) software program running on most major UNIX operating systems. It provides low-cost Internet connectivity.

MS PPP/SLIP software connects hosts and IP networks over links ranging in cost and performance from dial-up asynchronous modems through high-speed lines. The Morning Star PPP Version 1.4.1.7 adds support for the Solaris 2.5 operating system, along with most other major UNIX operating systems, and is available immediately through Progressive Systems' Sales Division and resellers.

Contact Progressive Systems, phone: (800) 558-7827 or (614) 326-4600, e-mail: sales@MorningStar.Com, <http://www.Progressive-Systems.com>.

Protect your investment in HP technology.

As a member of Interex, the International Association of Hewlett-Packard Computing Professionals, you can influence the decisionmaking process that determines the fate of your HP equipment. Through our many advocacy channels to Hewlett-Packard, our members actually become involved in the development process. These well established channels include:

- Interex annual surveys
- Conference roundtables
- Special Interest Groups

Make yourself heard!

Find out more about advocacy and all of the benefits of Interex membership. Just call one of our Membership Representatives at 800.INTEREX, 408.747.0227, or e-mail us at membership@interex.org.

Interex Online
<http://www.interex.org/>

interex

The International Association of Hewlett-Packard Computing Professionals

Make the UNIX to MPE Connection

IX/92™

Full featured HP terminal emulation

New! Version 6

**Faster File Transfer
NS/VT Network Option
Enhanced Script Language**

Available for:


HP-UX Interactive UNIX SCO UNIX SunOS/Solaris

Software Licensing Corp., Suite 280, 930 Tahoe Blvd. Unit #802
Incline Village, NV 89451-9436

Phone: (800) 831-0882 or (702) 832-0881 Fax: (702) 832-0883

All trademarks are the property of their respective holders.

CIRCLE 180 ON READER SERVICE CARD



WHY IS QEDIT™ SO FAMOUS?

- ✓ Qedit provides easy-to-use full-screen editing
- ✓ Qedit turbo charges PowerHouse and Cobol programmers
- ✓ Qedit's powerful commands put you in control
- ✓ Boosting programming productivity for twenty years



For more information find your local distributor listed below or in North America call 1-888-ROBELLE.

Qedit is available on HP 3000 and HP-UX computers.

Robelle
CONSULTING LTD.

Unit 201
15399-102A Ave.
Surrey, B.C.
Canada V3R 7K1

Toll free: 1-888-762-3553
Phone: (604) 582-1700
Fax: (604) 582-1799
E-mail: info@robelle.com
Web: www.robelle.com



Official Robelle Distributors

Australia, New Zealand.....61 2 484 3979	Greece, Italy.....30 1 777 0561	Scandinavia.....46 8 623 00 50
Central America.....502 2 314786	Holland, Belgium.....31 13 5215655	Singapore.....65 240 3432
Czech, Slovak Republics.....42 2 723 305	Hong Kong.....852 2609 1338	South Africa.....27 21 685 7809
France, Belgium.....33 1 69 86 60 00	Mexico.....52 5 813 1325	Switzerland, Austria.....41 31 981 06 66
Germany.....49 7621 689 190	Saudi Arabia, U.A.E.....966 1 477 4555	United Kingdom, Ireland...44 171 473 2558

Qedit is a trademark of Robelle Consulting Ltd.

CIRCLE 73 ON READER SERVICE CARD

SNA APIs

Microsoft Corporation and Parker Software Inc. have announced a native UNIX client for Microsoft SNA Server. Similar clients are available from Microsoft for Windows 3.x, 95, and NT Workstation and other operating systems. The SNA client allows LAN-based clients to connect through SNA Servers to IBM mainframe and AS/400 systems. The UNIX client will be made available for all popular UNIX variants.

The UNIX client relies entirely on Microsoft SNA Server, which runs on Microsoft Windows NT Server Version 3.51 or 4.0. The UNIX client dynami-

cally locates available SNA Server computers and downloads the appropriate client configuration automatically. The UNIX client benefits from advanced SNA Server features such as client-server data encryption, load balancing, and fault tolerance. The UNIX client for SNA Server requires much less memory, CPU, and system resources than a conventional UNIX-based SNA stack.

The single-unit price is expected to be \$1,495, with a 50-percent discount for competitive product upgrades.

Contact Parker Software, phone: (212) 843-0140, fax: (212) 889-5281, e-mail: sales@ParkerSoftware.com.

Web-based Document Management

Uniplex Software and Sybase, Inc. have announced a strategic alliance that will deliver fully recoverable Web Sites and enterprise-wide document management to users of Sybase SQL Server. Companies can get a free document management system for up to 15 users that is optimized for the Sybase SQL Server database. With this software, users can run a fully functioning pilot to manage documents across the enterprise, their intranet, or even their Web Site, without incurring any cost.

onGO DMS makes full use of Sybase database features, ensuring system integrity through a two-phase commit protocol between the document store and catalogue; 24-hour availability of the stored documents through online backup; and document accuracy through version and revision management.

Companies can register on the Uniplex Web Site, <http://www.uniplex.com>, for their free 15-user document management system.

Contact Uniplex Software, phone: (916) 985-3617, fax: (916) 985-6008.

Change Management Suite

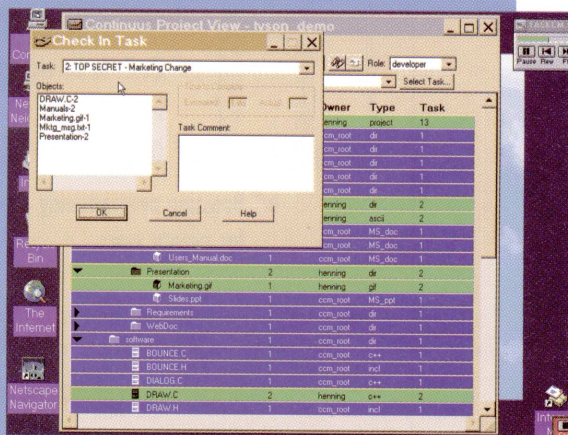
Continuous Software Corporation has shipped Version 4.2 of the Continuous Change Management Suite, which enlarges the scope of SCM beyond individual files to encompass the entire set of files involved in a given programming task.

Continuous' new Task-Based CM eliminates the need for developers to maintain lists of which files they've edited to complete each enhancement and bug fix. Task-Based CM also allows changes related to an enhancement or bug fix to be backed out of a software release with a mouse click.

Continuous Task-Based CM is easily modified to match the workflow of any desired software development process. Creating a new stage of workflow requires only the creation of task folders and a few additional mouse clicks.

The price of the server license for the Continuous Change Management Suite is \$9,995. Client pricing for the full suite is \$2,495 per seat for any UNIX or Windows platform. The software runs under HP-UX, Windows NT, and other operating systems.

Contact Continuous Software, (800) 820-1995 or (714) 453-2200, fax: (714) 453-2276.



Continuous, Change Management Suite

Remote Maintenance System

ARGOS SYSTEMES has announced the Ariane II telecommunications package, which features bi-directional file transfer and programming capabilities that enable companies to connect directly to customers' systems for real-time maintenance, troubleshooting, applications assistance, and system upgrades. The package operates over switched telephone networks or ISDN lines and readily handles interactive

communications with PCs running under different operating systems.

Transmission/reception capabilities include data encrypting, automatic error corrections, statistical reporting, transmission time estimation, line status analysis, log-in management, and real-time data compression/decompression.

Ariane II runs under HP-UX and other platforms.

ARGOS SYSTEMES welcomes inquiries from prospective partners interested in distributing its telecommunications products in the U.S.

Contact French Technology Press Office, phone: (312) 222-1235, fax: (312) 222-1237.

Enterprise Management

NOVADIGM, Inc. has announced a strategic alliance with BMC Software, Inc. to jointly develop for NOVADIGM's Enterprise Desktop Manager (EDM) new capabilities that optimize the management of BMC Software's PATROL management suite. During the initial phase of this alliance, a customized version of EDM that automatically distributes and manages changes to the PATROL environment will be provided to joint customers. Provided with Release 3.1 of PATROL, these new capabilities enable PATROL administrators to use EDM to update their production environment enterprise-wide without manual intervention.

The companies also announced that as a subsequent phase of the alliance, NOVADIGM's "desired-state" management technologies will be embedded in BMC Software's current and future Open Systems products.

Contact NOVADIGM, phone: (800) 662-6682 or (201) 512-1000, <http://www.novadigm.com>.

Full-screen Editor

Robelle Consulting Ltd. has announced Qedit/UX Version 4.5, which offers a new Screen mode that works with most VT terminals, such as VT100 and VT220. Unlike Visual mode, Screen mode does not rely on the block-mode feature of HP terminals. Screen mode allows users to make changes to the entire screen, just as they would on a PC editor. In addition to expanding the functions open to existing VT users, this simplified interface makes Screen mode a good introduction to the HP-UX operating system.

Users can now print to an HP-UX device or an attached printer and obtain more flexibility in working with "include" and "use" files.

All Qedit enhancements are compatible with HP-UX 10.0.

User can stay in Qedit to run programs, execute shell commands (Qedit knows which shell the user prefers), and shell scripts. Qedit also has a powerful command mode for editing files interactively or in shell scripts.

Contact Robelle, (888) ROBELLE or (604) 528-1700, fax: (604) 582-1799, e-mail: info@robelle.com.

Finite-Element Software

FRAMASOFT+CSI has announced the SYSTUS+ finite-element package, which includes an integrated automatic mesh generator, application-specific user interfaces, high-performance iterative solvers, a materials database, and solution post-processors. SYSTUS+ for workstations and supercomputers replaces two general-purpose packages, SYSTUS and MOSAIC.

A comprehensive set of modules handles a range of analysis problems in mechanical, thermal, electromechanical, production processes, and coupled analysis. Each of these modules is a dedicated tool for specific applications, operating as an integrated package with pre- and post-processors. Terminology appropriate for individual disciplines enables users to readily interact with the software.

SYSTUS+ has a Motif-based GUI with online hypertext help. Data may be exchanged with other software via standard protocols such as IGES and STEP. Access to PC-based spreadsheet data and word processing packages facilitates the preparation of analysis reports.

Contact French Technology Press Office, phone: (312) 222-1235, fax: (312) 222-1237.

Financial Institution Software

DIAGRAM has announced the Dome capital market and fund management software for banks, insurance companies, pension fund administrators, stock brokerage firms, mutual fund companies, and other financial institutions. Operating on either UNIX or Windows NT operating systems, the software is compatible with standard relational database management systems, including Sybase and Oracle.

Application modules include essential functions at three levels for capital markets and banking. A front office module covers pricing, time to market, and deal capture. Middle office functions encompass position tracking, profit and loss calculation, and risk management. Back office tasks include message processing, payment and delivery, routine accounting, and management and statutory reporting.

Software may be delivered as a turnkey system. A toolkit is also available for users to develop their own specific applications.

Contact French Technology Press Office, phone: (312) 222-1235, fax: (312) 222-1237.

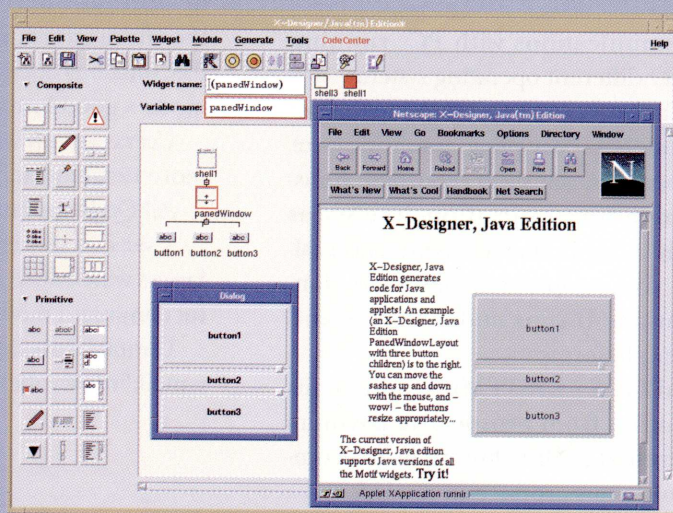
Motif to Java Migration

Imperial Software Technology has announced X-Designer: Java Edition, reportedly the only tool that gives developers a powerful migration path to Java for Motif interfaces. X-Designer can now generate code for Motif, Windows, and Java.

Advanced Motif application interfaces can be built with the knowledge that later deployment to any Java platform can be accomplished. Applications that are currently deployed internally may become Web-enabled client-server intranet applications. In addition, Java interfaces built with X-Designer: Java Edition may be ported back to Motif. Existing interfaces—including those not built using X-Designer—may be converted to the Java model.

X-Designer: Java Edition is priced at \$3,500 for the first license, including the first year's maintenance and upgrades. X-Designer: Java Edition runs under HP-UX and other UNIX platforms.

Contact Imperial Software Technology, phone: (415) 688-0200, fax: (415) 688-1054, e-mail: sales@ist.co.uk, <http://www.ist.co.uk>.



X-Designer: Java Edition

SAP-Certified Application Interface

Legato Systems, Inc. has announced that Legato's Application Interface for SAP's R/3 has been certified by SAP for use under Solaris, HP-UX, and AIX.

An implementation of SAP's BC-BRI specification, Legato's Application Interface for SAP's R/3, provides an interface between SAP's backup utilities and Legato's NetWorker storage management solution. When implemented, this system lets users perform backup and restore to NetWorker using SAP's *brbackup*, *brarchive*, *brrestore*, and *sapdba* utilities.

Legato's Application Interface for SAP R/3 supports cloning of savesets containing SAP backups, a capability that expedites creation of copies for off-site storage. The system can handle raw partition data in addition to file system data, and also features optional client-side compression.

Pricing for HP-UX is \$20,000 list per NetWorker server. Installation and consulting services are available.

Contact Legato Systems, phone: (415) 812-6000, fax: (415) 812-6032, <http://www.legato.com>.

File System Performance

EAGLE Software has announced DISK_PAK OnLine! designed to improve system performance without interrupting work or requiring the disk to be unmounted. DISK_PAK OnLine! allows system administrators to optimize their systems on the fly. It is currently available for HP-UX and other UNIX platforms.

DISK_PAK OnLine! safely eliminates file system fragmentation and clusters frequently accessed files for peak file system responsiveness. All pieces of each file are grouped together, eliminating access fragmentation. Space fragmentation is also handled by grouping available space for quicker disk access time.

DISK_PAK OnLine! is transparent to the user, running in the background to keep the system performing at optimum levels without any downtime. The enhanced version also defragments the root file system.

By making copies of individual file elements during reorganization, DISK_PAK Online! keeps the file system consistent and error free, even if the power fails during a reorganization.

Contact EAGLE Software, phone: (800) 477-5432 or (913) 823-7257, or e-mail: info@eaglesoft.com.

New from HP

HP OpenView Oracle7

Oracle Corporation has announced a licensing agreement that gives HP the right to resell the Oracle7 relational database for the HP OpenView network and systems management framework. Oracle7 for HP OpenView allows users to employ a relational database for storing, managing, and accessing systems management data. Oracle7 allows HP OpenView users to perform advanced queries or reports, ensure data integrity through row-level locking, and maintain transaction control throughout the network.

The new Oracle7-based repository is available now and will be sold through

HP's direct sales force and partners. HP and Oracle's support organizations have entered into an agreement to share information to jointly solve customer support issues. Like OpenView applications, the repository will be priced per installation. Oracle will provide technical support to HP's Response Centers on a global basis.

HP is a member of the Systems Management Tools Initiative (SMTI), an Oracle-sponsored program that fosters closer working relationships between Oracle and providers of complementary administration tools.

Internet Portfolio

HP has announced HP Domain solutions and new Internet offerings to help large and small organizations get control of their business on the Net. These solutions span Microsoft Windows NT and HP-UX operating environments and include the following:

- HP Domain Partner Program—a program designed to foster the development of Internet- and intranet-based computing solutions across operating environments (for example, Windows NT and UNIX systems)
- HP Domain Business Suite—reportedly the first business-ready Internet and intranet software suite to incorporate unparalleled system, network, and application management based on HP OpenView, Netscape SuiteSpot server software, and integrated messaging from HP OpenMail
- HP Domain Enterprise Server family—UNIX Internet and intranet servers bundled with HP Domain Business Suite
- HP NetServer Web Master—an integrated Internet PC server solution

based on Microsoft Windows NT that gives small to medium-sized businesses Internet access and presence in less than an hour

- channel programs for HP resellers
- enhanced consulting services and support programs

The HP Domain Partner Program consists of three core elements: a broad range of co-marketing efforts, discounts on HP hardware and software for porting and development, and access to a wide range of financing options.

HP's HP-UX-based Domain servers run on HP's 64-bit PA-8000 RISC processors and are preconfigured with the HP Domain Business Suite, which combines Internet software from Netscape with HP management and messaging solutions that extend the capabilities of the Internet for information access, collaboration, and core business transactions. They are easy to use and ready to go and include the following key components: management based on HP OpenView, Netscape SuiteSpot software, and HP OpenMail.

The HP NetServer Web Master is said to be the industry's first integrated Windows NT-based PC server solution that gives small to medium-sized businesses Internet access and presence in less than an hour. Optimized and integrated with Internet hardware and software from partners such as Intel and Microsoft, Web Master arrives preconfigured by HP and its Channel Partners to help businesses get online fast. Specialized technical expertise or staff is

Enterprise Production Management

New Dimension Software has introduced a new version of CONTROL-M—Option for R/3, an interface module allowing the company's CONTROL-M enterprise production control and scheduling software to manage the scheduling and monitoring of SAP R/3 workloads.

An integral part of the company's Enterprise Production Management solution, CONTROL-M allows operations personnel to easily manage and automate the setup, scheduling, and execution of processes running across a multiplatform computing environment.

CONTROL-M—Option for R/3 allows users to automate the scheduling and management of both ABAP/4 reports and SAP jobs. CONTROL-M—Option for R/3 is also enhanced with features for defining and running SAP jobs directly from CONTROL-M.

CONTROL-M—Option for R/3 is priced starting at \$32,000.

Contact New Dimension Software, (800) 347-4694, ext. 522 (North America), 972-3-645-1111 (Israel), fax: (714) 756-3900, <http://www.ndsoft.com>.

not needed to install and configure the hardware and software.

The Worldwide Customer Support Organization also introduced a new service offering called the HP NetServer Web Master Support Pak. The new support offering gives small and medium-sized businesses the option to receive service for their HP NetServer Web Master either directly from their dealer or from the HP channel. One five-incident support pak is included with the solution, and add-on support paks will be available and sold through local resellers.

SoftBench Project: GV

HP has announced SoftBench Project: GV, a project and process access tool that combines Web technology with configuration-management technology

to give software-development processes increased visibility.

Project: GV's special integration of a powerful Web-accessible configuration-management system with proven process methodology helps organizations reduce development-cycle time, gain greater project visibility, achieve better control of project documents, make project documents available on the Web, and document the project development life cycle online. Using Project: GV technology, some HP developers have reduced their development-cycle times by about 50 percent, HP notes.

The SoftBench CM component of Project: GV provides the means for versioning, controlling, and managing life-cycle templates, process, and project documents. Project: GV integrates the SoftBench CM server and the organization's Web server, thereby making files in the SoftBench CM archive available, with versioning information, to Web browsers such as Netscape Navigator and Microsoft Explorer.

Project: GV includes project and process template sets based on IEEE and Software Engineering Institute (SEI) standards for software-process effectiveness. These templates provide guidance for building a solid process and life cycle online that can be customized to suit the user's own process needs.

The key to Project: GV's power is the SoftBench CM link to the Web, which connects individual users and their Web server to the SoftBench CM configuration-management server. The SoftBench CM server manages the versioning of life cycle and project archives, including source code. Without the SoftBench CM link to the Web, none of the version information is accessible in a browser. With their Web server linked to Project:

GV's SoftBench CM server, the organization can gain control over the project archives and still allow them to be viewed through a Web browser, checked out, and worked on.

Companies that develop their software using geographically dispersed resources will find Project: GV particularly advantageous.

Critical Systems Support

HP has introduced HP Critical Systems Support, a suite of support capabilities designed to ensure minimal productivity loss in enterprise-class distributed-computing environments. The new service expands HP's range of high-availability services, furthering the company's lead in providing support alternatives to keep business-critical systems running at optimum levels. Under HP Critical Systems Support, HP is committed to resolve all HP system hardware problems within six hours of receiving a customer repair call.

The core service associated with HP Critical Systems Support includes an assigned account team experienced in high-availability hardware, software, and network support. When a problem occurs, customers are immediately connected to a system-recovery team with access to HP's critical-parts network and other resources that may be required to repair system hardware within six hours. In addition, phone-in software assistance is available 24 hours a day, 365 days a year. Another core service feature is quarterly "patch" management reviews designed to ensure that appropriate patches have been installed on supported UNIX systems.

Users can add enhanced on-site and remote software support, as well as house critical parts at their own sites or

at an HP site.


HP Critical Systems Support is available now for HP-UX enterprise environments, and HP expects to expand the service to Microsoft Windows NT environments by the second half of 1997. Pricing ranges from \$100,000 to \$120,000 annually, depending on server size. ■

Attention vendors: New product announcements should be sent to New Products Editor, hp-ux/usr magazine, Interex, P.O. Box 3439, Sunnyvale, California 94088-3439, USA, or e-mail: pollace@interex.org.

Deadline for submission is two months prior to publication.

HP[®] SALES • RENTALS • SERVICE

1-800-422-4872



WORKSTATIONS

B Series
C Series
J Series
Visualize Graphics

**COMPLETE LINE 700 & 800
SERIES IN
STOCK!**

BUSINESS SERVERS


E Series
D Series
G - H - I Series
K Series

AND MORE...

Add On Memory
Disc Arrays
Opticals
Netservers

Test & Measurement
Personal Computers
Data Acquisition
Notebooks

X Terminals
Printers
Plotters
Disc Drives



THE RIGHT EQUIPMENT. RIGHT NOW.

TSA

2040 West Sam Houston Parkway N. • Houston, Texas 77043

713/935-1500 • Fax: 713/935-1555 • <http://www.tsa.com> • E-mail: info@tsa.com

New kids on the block...

“ I bet there's a whole bunch of ads in this magazine for **tape libraries**. But you won't want to miss what my dad has to say, 'cuz his SL-400's really neat—it's easy to use, it won't break, and you can even buy it from him, too. ”



CIRCLE 103 ON READER SERVICE CARD

Consider the advantages of the **INTEREX DIRECT MAIL SERVICE**

- ◆ 46,000 Hewlett-Packard computer users in our database
- ◆ A one-week turnaround time
- ◆ Less than 2% returns because of incorrect addresses
- ◆ 74% of our users are their company's decision makers
- ◆ Full refund for all nixes (bad addresses)
- ◆ 15% discount on all orders over 10,000 labels
- ◆ List updated weekly

We offer the most value for your money through direct mailings and the INTEREX Card Deck.

Contact Phil Nguyen, INTEREX,
P.O. Box 3439, Sunnyvale,
CA 94088-3439, USA

(800) INTEREX • (408) 743.4619 FAX; (408) 747.0947

...with seven years of experience

and a customer list that includes IBM, Disney, Shell, Chrysler, Price Waterhouse, Sallie Mae, and others.

Here's our story. For two decades, **IGM Communications** made tape access devices for the broadcasting industry. With great design and leading-edge technology, IGM captured over 80% of its market.

Seven years ago, IGM transferred its expertise in tape handling to the computer data storage market. IGM's ATL-5000 was the industry's first 8mm tape library. It's still being used in a wide range of applications—from storing animated drawings to collecting seismic

data—winning customer loyalty for its exceptional reliability.

In 1994, I bought IGM, launched a new division, **Straightline**, and began an aggressive strategy to produce robotic tape libraries. Our unique solution to data backup combines **4mm DAT technology** with easy operation, proven dependability, and manufacturer-direct sales and support.

So you see, we might be the new kids on *this* block, but we've been living in your neighborhood for years.



Get to know us better
by calling

1.800.458.1273



straightline

677 120th Ave. NE, Suite 146 • Bellevue, WA 98005
www.igm.com • e-mail: igm@igm.com

Software Compatibility:

UNIX Environments: Legato Networker, IBM Adstar, Raxco, Spectralagic, OSM, Tivoli; Netware Environments: Cheyenne ARCserve, Palindrome; NT Environments: Arcada Backup Exec

©1996 Straightline. All rights reserved. Straightline is a division of IGM Communications. All brand and product names are property of their respective holders.

SIGs

If your area of expertise involves HP systems, there is a place you can go to get the kind of specialized, top level information that's valuable to you. Interex Special Interest Groups (SIGs).

As a significant part of Interex's advocacy process, SIGs provide members with an opportunity to increase their knowledge through one-on-one interaction with HP representatives and their computing colleagues.

Find out more by calling us at 1.800.INTEREX (408.747.0227) or e-mailing us at membership@interex.org.

interex

Interex Online
<http://www.interex.org/>

The International Association of Hewlett-Packard Computing Professionals

Focused

on the

Future

May & Speh provides a broad range of technological expertise and services in data management and data center outsourcing. Our growth has been steady and sure. In 1992, our sales were \$32.2 million; just completing FY 1996, sales were \$77.2 million. While other companies are downsizing, May & Speh has been constantly adding to its professional workforce. Our aggressive business strategies and the continuous enhancement of our state-of-the-art technology keep us focused on the future. Committed to our clients and our talented employees, we will continue to be the industry's dominant force well into the 21st Century.

UNIX SYSTEMS ADMINISTRATORS

You will support our two terabyte+ multiplatform environments, install UNIX software and hardware, and ensure the optimal performance and integrity of our UNIX systems. Qualified candidates possess a minimum of 4 years UNIX systems experience, are familiar with various flavors of UNIX-IBM, HP, SGI, and have experience setting up and configuring a networked UNIX environment. System performance monitoring and tuning experience is desirable. A professional demeanor and excellent communication skills are a must.

Focus on your future at **May & Speh**. We offer highly competitive salaries and a comprehensive benefits package including Employee Stock Ownership Plan, tuition reimbursement and PC purchase program. Please send your resume with salary history and expectations to: **May & Speh, Attn: Dept. HR-HP1/97, 1501 Opus Place, Downers Grove, IL 60515-5713; fax: (630) 719-0290; e-mail: mark.jeske@mayspeh.com**

Equal Opportunity Employer M/F/D/V.



MAY & SPEH
Technology, Talent, Service.

Advertiser's Index

Please call or fill out adjacent card for further product information.

READER SERVICE NUMBER	ADVERTISER	PAGE #
3	AMC Computer Services (508) 670-9395/Fax (508) 670-9327	35
23	BMC Software (800) 841-2031	2,3
24	Centon Electronics (714) 855-9111/Fax (714) 855-1144	C-4
6	CLAM Associates (617) 621-2542 or e-mail: marketing@clam.com	55
153,154	Computer Solutions (407) 649-0123/Fax (407) 649-1407	21,51
28	Concorde Technologies (800) 359-0282/Fax: (619) 536-5500	7
37	Confluent (800) 780-2838 or email: info@confluent.com	77
140	C.S.U. Industries, Inc. (516) 239-4310/Fax: (516) 239-8374	25
76	Dallastone (603) 647-8168/Fax: (603) 624-2466	29
44	Design 3000 Plus, Inc. (503) 585-0512/ Fax: (503) 585-1706	39
20	Digital StorageWorks (800) 663-1966 or (603) 890-6700	17
16	IEM, Inc. (970) 221-3005 or (800) 321-4671/e-mail: info@iem.com	13
171	I/O Data Systems, Inc. (216) 835-2211/Fax: (216) 835-0220	29
129	Lund Performance Solutions (541) 926-3800	59
85	May & Speh Fax: (630) 719-0290 or e-mail: Mark.jeske@mayspech.com	87
7	Monterey Bay Communications (408) 429-6144/Fax: (408) 429-1918	1

READER SERVICE NUMBER	ADVERTISER	PAGE #
74	Minisoft, Inc. (800) 682-0200/Fax: (360) 568-2923	63
72	ORBIT Software (800) 6-Online or (510) 837-4143	27
91	Personal Productivity Tools (708) 620-5000/Fax: (708) 691-0718	69
73	Robelle Consulting Ltd. (604) 582-1700/Fax: (604) 582-1799	79
138	SAIC (800) 447-4373/Fax: (619) 552-5253	C-2
180	Software Licensing Corp. (702) 832-0881	79
165,103	Straightline (800) 458-1273/Fax: (206) 865-8314	71,86,87
175	Syncsort (201) 930-8200/Fax: (201) 930-8290	9
163	Syntax (206) 838-2626	5
123	Taurus Software (415) 961-1323, ext. 100	C-3
78	Technical & Scientific Application (800) 422-4872/Fax: (713) 935-1555	85
49	Ted Dasher & Associates (800) 638-4833/Fax: (205) 591-1108	61
40	VESOFT, Inc. (310) 282-0420/ Fax: (310) 785-9566	77
127	VIA Technology (800) 842-8324	61
132	Xess Corp. (800) 961-7840/Fax: (919) 387-1302	75

Data Movement Solved!

HP ORACLE UNIX

IMAGE MIGRATION

FLAT FILES CODASYL DBMS DEC VMS TEST DATABASES
ALLBASE ARCHIVING DECISION SUPPORT

Thanks, Warehouse.

Warehouse does two things. Move and manipulate data. Almost anywhere. No programming needed.

If you need to cross HP, UNIX or VAX platforms. Or move between Image, Oracle, DBMS and other databases or applications. No problem. If you need to migrate legacy data across the network straight into relational databases. Done. It arrives as native data.

Data warehousing is just as easy. In minutes, you can archive and retrieve selected data from any media. Make

test and decision support databases. Think of the advantages. No downtime. No big intermediate flat files. And total platform transparency.

Hundreds of companies depend on Warehouse. And Taurus is a certified "HP User Reference" company as well as a participant in the "Oracle Business Alliance Program." So if you move data, now there's a real solution. Warehouse from Taurus. Call today for details.

Because Warehouse moves data.

Taurus Software

1032 Elwell Court, Palo Alto, CA 94303

Phone: 415-961-1323 x100 • Fax: 415-961-1454 • E-mail: sales@taurus.com • Web: www.taurus.com

CIRCLE 123 ON READER SERVICE CARD



Before You Travel the Crowded Road of Memory Upgrades, Consider This:

- Centon's experience: we've been in business for over 18 years.
- Centon's selection: we carry over 1,500 different modules.
- Centon's warranty: all products come with a lifetime guarantee.
- Centon's commitment: we're an authorized Apple Developer, Sun Licensee and a partner in IBM's warranty program.

So why not bypass the other memory suppliers? Take the *Memory Expressway* to unparalleled service, aggressive pricing and guaranteed product.

After all, for memory upgrades, all roads lead to Centon Electronics.

Pathfinder Plus™


Electronic Memory Guide



CENTON
ELECTRONICS, INC.

Call Centon today to get your free copy of our Electronic Memory Guide disk: "Pathfinder Plus."

Centon and the Centon logo are registered trademarks of Centon Electronics, Inc. All other brand names and trademarks are property of their respective owners.
©1996 Centon Electronics, Inc.



Pictured is a Centon manufactured memory module for the Apollo 9000 Series. Call Centon today for pricing and memory configurations for other Hewlett-Packard compatible upgrades.

1-800-836-1986

CENTON®
ELECTRONICS, INC.

Centon Electronics, Inc. Irvine, CA (714) 855-9111 Fax (714) 855-6035 <http://www.centon.com>

Proudly distributed by:

1 (800) 255-4489 AmeriQuest
1 (800) 223-7081 AmeriQuest/Robec

AMERIQUEST
COMPUTER 2000 GROUP

**INGRAM
MICRO**

"Leading The Way In Worldwide Distribution"
1 (800) 456-8000

CIRCLE 24 ON READER SERVICE CARD